arXiv:1702.06234v1 [stat.ML] 21 Feb 2017

On a Class of First-order Primal-Dual Algorithms for Composite Convex Minimization Problems

Seyoon Ko[1], Donghyeon Yu[2], and Joong-Ho Won[*1]

[1]Department of Statistics, Seoul National University
[2]Department of Statistics, Inha University

Abstract

Many statistical learning problems can be posed as minimization of sum of two convex functions, one typically non-smooth. Popular algorithms for solving such problems, e.g., ADMM, often involve non-trivial optimization subproblems or smoothing approximation. We study two classes of algorithms that do not incur these difficulties, and unify them from a perspective of monotone operator theory. The result is a class of preconditioned forward-backward algorithms with a novel family of preconditioners. We analyze convergence of the whole class of algorithms, and obtain their rates of convergence for the range of algorithm parameters where convergence is known but rates have been missing. We demonstrate the scalability of our algorithm class with a distributed implementation.

***Keywords*—** non-smooth optimization, convex optimization, operator splitting, generalized lasso, group lasso, distributed computing, multi-GPU

1 Introduction

Many statistical learning problems can be formulated as an optimization problem of the form

$$\text{minimize } f(x) + h(Kx), \tag{1}$$

where $x \in \mathbb{R}^p$ and $K \in \mathbb{R}^{l \times p}$, and both f and h are closed, proper, and convex. In this article we consider the case when f is L_f-strongly smooth,[†] but h is not necessarily smooth. The minimization problem (1) can be equivalently written as the following primal-dual form

$$\min_{x \in \mathbb{R}^p} \sup_{y \in \mathbb{R}^l} f(x) + \langle Kx, y \rangle - h^*(y), \tag{2}$$

[*]wonj@stats.snu.ac.kr

[†]A function f is L_f-strongly smooth if the gradient ∇f of f exists and is L_f-Lipschitz continuous, i.e., $\|\nabla f(x) - \nabla f(x')\|_2 \le L_f \|x - x'\|_2$, where $\|x\|_q = \left(\sum_{i=1}^p |x_i|^q\right)^{1/q}$.

where $h^*(v) = \sup_{u \in \mathbb{R}^l} \langle u, v \rangle - h(u)$ is the convex conjugate of h.

A pinnacle instance of (1) is high-dimensional penalized regression with structured sparsity assumptions, which solves the following optimization problem.

$$\text{minimize} \quad \sum_{i=1}^{n} l_i(a_i^T x, b_i) + H(Dx), \tag{3}$$

where $x \in \mathbb{R}^p$ is the vector of regression coefficients, $D \in \mathbb{R}^{l \times p}$ is the structure-inducing matrix, $l_i : \mathbb{R}^2 \to \mathbb{R}$ is the loss function, and H is the penalty function, which is typically non-smooth. The set $\{(a_i, b_i) : i = 1, \ldots, n\}$ constitutes a training sample. Note (1) includes (3) as a special case, by identifying $f(x) = \sum_{i=1}^{n} l_i(a_i^T x; b_i)$, $H(u) = h(u)$, and $K = D$. Strongly smooth loss functions arise in many important problems, e.g. $L_f = \|A^T A\|_2$ for the quadratic loss $l_i(\cdot, b_i) = \frac{1}{2}(\cdot - b_i)^2$, and $L_f = \frac{1}{4}\|A^T A\|_2$ for the logit loss $l_i(\cdot, b_i) = -b_i(\cdot) - \log(1 + e^{-(\cdot)})$; $A = [a_1, \ldots, a_n]^T$ is the data matrix, and $\|A^T A\|_2 = \sup_{x \neq 0} \frac{\|A^T A x\|_2}{\|x\|_2}$. For an example of nonsmooth penalty functions, the generalized lasso (Tibshirani and Taylor, 2011) uses $H(z) = \lambda \|z\|_1$ for $z \in \mathbb{R}^l$ and $\lambda > 0$. In particular, choosing D as a finite difference matrix on a graph yields the fused lasso (Tibshirani et al., 2005).

Problem (1) is challenging in general because the non-separability of the non-smooth part hampers use of efficient methods such as element-wise soft-thresholding (Friedman et al., 2007; Beck and Teboulle, 2009) available when K is the identity matrix and h is separable. Much efforts have been made to solve (1) in large scale efficiently. Many existing approaches require either smooth approximation of the non-smooth penalty function, solving quadratic minimization subproblems, or evaluating a non-trivial proximity operator of $h \circ K$ for each iteration (see Section 2 for proximity operators). The former introduces an additional parameter that determines the degree of smoothing, which is difficult to choose in practice; the latter two may become computationally burdensome for large-scale problems. See Appendix A for a brief summary. Formally the same problem (1) has been also studied in the image processing literature (Rudin et al., 1992; Goldstein and Osher, 2009; Chambolle and Pock, 2011). However, in statistical learning applications where data matrix is usually large and unstructured, avoiding quadratic minimization is in a greater need.

We study two classes of primal-dual algorithms for solving (1) that require neither smoothing nor quadratic minimization. The proximity operator employed by these algorithms depends only on h but not K. Thus they are simple to implement and attractive for parallel and distributed computation. Non-smooth loss functions, such as the hinge loss, are also allowed in (3) by a different choice of f, h, and K in (1). Therefore these algorithms are versatile in handling many important statistical learning problems.

Our contributions can be summarized as follows. First, we connect the two classes of algorithms, which have been regarded as separate classes in the literature, from a perspective of monotone operator theory and show that they are essentially the same

class of algorithms: they can be posed as a preconditioned forward-backward splitting algorithm (see, e.g., Combettes and Wajs, 2005) sharing the *same* preconditioner. From this connection we propose a new, broader family of preconditioners and the corresponding class of forward-backward algorithms. Second, we analyze convergence of this new class of algorithms in a unified fashion. The cited connection allows us to derive conditions for convergence for the whole class of algorithms, and obtain its rates of convergence. For both classes of algorithms the rates have been available under much more stringent conditions than that for convergence; we close this gap. Third, by exploiting the advantages of the studied algorithms, we provide a distributed algorithm that split features in multiple devices when the data do not fit in the memory of a single device.

Notation. We restrict our attention to finite-dimensional Euclidean spaces. We denote the standard inner product $x^T y$ by $\langle x, y \rangle$, and the associated norm by $\|\cdot\|_2$. The spectral norm, or the maximum singular value, of matrix A is also denoted by $\|A\|_2$. That a symmetric matrix M is positive (semi)definite is denoted by $M \succ 0$ ($M \succeq 0$); $A \succ B$ refers to $A - B \succ 0$, etc. For $M \succ 0$, we define its associated inner product and norm by $\langle x, x' \rangle_M = \langle Mx, x' \rangle$ and $\|x\|_M = \sqrt{\langle x, x \rangle_M}$, respectively. For a symmetric matrix A, $\lambda_{\max}(A)$ and $\lambda_{\min}(A)$ respectively denote the maximum and minimum eigenvalues.

2 Monotone operator theory

Here we briefly state necessary results from monotone operator theory. For more details, see Bauschke and Combettes (2011).

Set-valued operators. A set-valued operator $T : \mathbb{R}^n \to 2^{\mathbb{R}^n}$ maps a vector $z \in \mathbb{R}^n$ to a set $T(z) \subset \mathbb{R}^n$. The graph of T is denoted by $\mathbf{gra}\, T = \{(z, w) \in \mathbb{R}^n \times \mathbb{R}^n : w \in T(z)\}$. When $T(z)$ is single-valued, i.e., $T(z) = \{w\}$, T is a function, and we write simply as $T(z) = w$. We use I to denote the identity operator, i.e, $I(z) = z$. When no confusion incurs, we also use Tz to mean $T(z)$. In particular, when T is a single-valued linear operator, Tz is identified with a multiplication of the corresponding matrix $T \in \mathbb{R}^{n \times n}$ by a vector z. The set of zeros of T is defined as $\mathbf{zer}\, T = \{z \in \mathbb{R}^n : 0 \in Tz\}$. The inverse of T is $T^{-1} : \mathbb{R}^n \to 2^{\mathbb{R}^n}$ such that $T^{-1}(w) = \{z \in \mathbb{R}^n : w \in Tz\}$, hence $\mathbf{gra}\, T^{-1} = \{(w, z) \in \mathbb{R}^n \times \mathbb{R}^n : w \in Tz\}$. The resolvent of T is $R_T = (I + T)^{-1}$. Scaling of an operator T by $t \in \mathbb{R}$ is defined by $(tT)(z) = tT(z)$. Composition of two set-valued operators $T_1 : \mathbb{R}^n \to 2^{\mathbb{R}^n}$ and $T_2 : \mathbb{R}^n \to 2^{\mathbb{R}^n}$ is defined by $T_2 T_1 z = \bigcup_{w \in T_1 z} T_2 w$.

Fixed points. An operator $T : \mathbb{R}^n \to 2^{\mathbb{R}^n}$ is called nonexpansive if $\|Tz - Tz'\|_2 \leq \|z - z'\|_2$ for all $z, z' \in \mathbb{R}^n$; it is called contractive if the inequality is strict. Any nonexpansive operator is single-valued. The set of fixed points of a single-valued operator T is denoted by $\mathbf{Fix}\, T$, i.e., $\mathbf{Fix}\, T = \{z : z = Tz\}$. For a contractive operator T, the fixed point iteration $z^{k+1} = Tz^k$ converges to a point in $\mathbf{Fix}\, T$, if $\mathbf{Fix}\, T \neq \emptyset$.

Averaged operators. An operator T is called α-averaged, $0 < \alpha < 1$, if $T = (1-\alpha)I + \alpha R$ for some nonexpansive operator R. Usually R is defined implicitly. Note that T itself is nonexpansive, and $\mathbf{Fix}\, T = \mathbf{Fix}\, R$. If T_1 is α_1-averaged and T_2 is α_2-averaged, then $T_1 T_2$ is α-averaged where $\alpha = (\alpha_1 + \alpha_2 - 2\alpha_1\alpha_2)/(1-\alpha_1\alpha_2)$. An α-averaged operator T is nonexpansive but not necessarily contractive, hence the fixed point iteration $z^{k+1} = Tz^k$ above may not converge to a fixed point even if $\mathbf{Fix}\, T \neq \emptyset$. In this case, the Krasnosel'skiĭ-Mann (KM) iteration $z^{k+1} = z^k + \rho_k(Tz^k - z^k)$ with a sequence $\{\rho_k\} \subset [0, 1/\alpha]$ such that $\sum_{k=0}^{\infty} \rho_k(1-\alpha\rho_k) = \infty$ ensures convergence.

Monotone operators. An operator T is called monotone if $\langle z - z', w - w' \rangle \geq 0$ for all $z, z' \in \mathbb{R}^n$ and for all $w \in Tz$, $w' \in Tz'$, and maximally monotone if it is monotone and there is no monotone operator T' such that $T \neq T'$ and $\mathbf{gra}\, T \subset \mathbf{gra}\, T'$. The resolvent of a maximally monotone operator is single-valued; it is $1/2$-averaged.

Cocoercive operators. A single-valued operator T is called γ-cocoercive if for some $\gamma > 0$, $\langle z - z', Tz - Tz' \rangle \geq \gamma \|Tz - Tz'\|_2^2$. A cocoercive operator is maximally monotone. If an operator T is γ-cocoercive with $\gamma > 1/2$, then $I - tT$ ($t > 0$) is $t/(2\gamma)$-averaged.

Subdifferential. An important example of a maximally monotone opeartor is the subdifferential of a convex closed proper function. A vector $g \in \mathbb{R}^n$ is a subgradient of a convex function ϕ at z if $\phi(z') \geq \phi(z) + \langle g, z' - z \rangle$, $\forall z' \in \mathbb{R}^n$. The subdifferential of ϕ at z is the set of subgradients at z: $\partial\phi(z) = \{g \in \mathbb{R}^n : \phi(z') \geq \phi(z) + \langle g, z' - z \rangle,\ \forall z' \in \mathbb{R}^n\}$. When ϕ is differentiable, $\partial\phi(z) = \{\nabla\phi(z)\}$. If ϕ is in addition closed and proper, $(\partial\phi)^{-1} = \partial\phi^*$ holds, where ϕ^* is convex conjugate defined by $\phi^*(w) = \sup_{z \in \mathbb{R}^n}\{\langle z, w \rangle - \phi(z)\}$. The resolvent of a maximally monotone subdifferential operator is the proximity operator: $R_{\partial\phi} = (I + \partial\phi)^{-1}(z) = \mathbf{prox}_{\phi}(z) = \arg\min_{z' \in \mathbb{R}^n} \phi(z') + \frac{1}{2}\|z' - z\|_2^2$.

Skew-symmetric operators. Another example of a maximally monotone operator is a skew-symmetric matrix. The sum of a maximally monotone operator and a skew-symmetric matrix is also maximally monotone.

Change of metric. Note that the notion of nonexpansiveness, averagedness, cocoercivity, and monotonicity of an operator requires the inner product $\langle \cdot, \cdot \rangle$ and its associated norm $\|\cdot\|_2$. We can appropriately define these concepts with respect to another inner product and its associated norm as well, say $\langle \cdot, \cdot \rangle_M$ and $\|\cdot\|_M$, for M a symmetric, positive definite matrix. In particular, averagedness of composition, convergence of the KM iteration, and averagedness of $I - tT$ for cocoercive T hold by substituting the inner products and norms by $\langle \cdot, \cdot \rangle_M$ and $\|\cdot\|_M$, respectively.

Forward-backward splitting. Some optimization problems can be translated to finding an element of $\mathbf{zer}\, T$ for an appropriate choice of maximally monotone operator T. Often T can be split into a sum of two maximally monotone operators F and G. If G is single-valued, then we can see

$$\begin{aligned} 0 \in T(z) &\iff (I + tF)(z) \ni (I - tG)(z) \\ &\iff z = R_{tF}(I - tG)(z), \end{aligned} \tag{4}$$

for $t > 0$. Equivalence (4) shows that $\mathbf{zer}\,(F+G) = \mathbf{Fix}(R_{tF}(I - tG))$, thus we may solve the problem of finding a zero of T by the following fixed-point iteration

$$z^{k+1} = R_{tF}(I - tG)(z^k). \tag{5}$$

This method is called the *forward-backward splitting algorithm*, and converges if G is cocoercive :

Proposition 1 (Bauschke and Combettes (2011), proof of Theorem 25.8)**.** *Let $F : \mathbb{R}^n \to 2^{\mathbb{R}^n}$ be maximally monotone, $G : \mathbb{R}^n \to \mathbb{R}^n$ be γ-cocoercive, and $\gamma > 0$. Let $t \in (0, 2\gamma)$ and $\delta = 2 - t/(2\gamma)$. Furthermore, let $\{\rho_k\}$ be a sequence in $[0, \delta]$ and let $z^0 \in \mathbb{R}^n$. Suppose $\mathbf{zer}(F+G) \neq \emptyset$ and set*

$$z^{k+1} = (1 - \rho_k)z^k + \rho_k R_{tF}(I - tG)(z^k). \tag{6}$$

Then the following hold:

1. $\|z^{k+1} - z\|_2^2 \leq \|z^k - z\|_2^2 \;\; \forall z \in \mathbf{zer}(F+G)$.
2. $\sum_{k=0}^{\infty} \frac{\delta - \rho_k}{\rho_k} \|z^{k+1} - z^k\|_2^2 \leq \|z^0 - z\|_2^2 \;\; \forall z \in \mathbf{zer}(F+G)$.
3. $\|z^{k+1} - z^k\|_2 \to 0$.

In addtion, if $\sum_{k=0}^{\infty} \rho_k(\delta - \rho_k) = +\infty$,

4. *$z^k \to z^\star$ for some $z^\star \in \mathbf{zer}\,(F+G)$.*

Note that $R_{tF}(I - tG)$ is a $1/\delta$-averaged operator, because R_{tF} is $1/2$-averaged, $I - tG$ is $t/(2\gamma)$-averaged. Thus (6) is a KM iteration. Choosing the relaxation coefficients $\rho_k \equiv 1$ reduces (6) to (5).

Preconditioning. In the derivation of (4), we can observe that the identity matrices in the first line can be replaced by an invertible matrix M, yielding a *preconditioned* forward-backward splitting algorithm

$$z^{k+1} = R_{tM^{-1}F}(I - tM^{-1}G)(z^k) \tag{7}$$

and its *relaxed* variant

$$z^{k+1} = (1 - \rho_k)z^k + \rho_k R_{tM^{-1}F}(I - tM^{-1}G)(z^k) \tag{8}$$

Preconditioning is useful when evaluating the resolvent $R_{tM^{-1}F}$ is easier than R_{tF}. If $M^{-1}F$ is maximally monotone and $M^{-1}G$ is cocoercive, Proposition 1 applies. It can be shown that if M is symmetric positive definite, $M^{-1}F$ is maximally monotone with respect to $\langle \cdot, \cdot \rangle_M$ (Combettes and Vũ, 2014), and $M^{-1}G$ is $\gamma\lambda_{\min}(M)$-cocoercive with respect to $\|\cdot\|_M$ (Davis, 2015). Therefore Proposition 1 holds for (7) and (8) with replacing $\|\cdot\|_2$ by $\|\cdot\|_M$, and γ by $\gamma\lambda_{\min}(M)$.

3 Two algorithm classes

We introduce the two classes of algorithms to study. We can write the first-order optimality condition for (1) as

$$\begin{aligned} 0 &= \nabla f(x^\star) + K^T y^\star \\ y^\star &\in \partial h(Kx^\star) \Leftrightarrow Kx^\star \in \partial h^*(y^\star), \end{aligned} \tag{9}$$

which is equivalent to the following fixed-point equation:

$$\begin{aligned} x^\star &= x^\star - \tau(\nabla f(x^\star) + K^T y^\star) \\ y^\star &= \mathbf{prox}_{\sigma h^*}(y^\star + \sigma K x^\star), \end{aligned} \tag{10}$$

for $\tau, \sigma > 0$. Note (10) is equivalent to the saddle-point condition for (2):

$$\mathcal{L}(x^\star, y) \leq \mathcal{L}(x^\star, y^\star) \leq \mathcal{L}(x, y^\star), \quad \forall x \in \mathbb{R}^p, \ \forall y \in \mathbb{R}^l,$$

where the saddle function $\mathcal{L}(x, y) = f(x) + \langle Kx, y\rangle - h^*(y)$ is convex in x and concave in y. Thus y can be interpreted as a dual variable.

The algorithms of our study solve (10) iteratively. The first class is due to Loris and Verhoeven (2011) and Chen et al. (2013):

$$\begin{aligned} \tilde{x}^{k+1} &= x^k - \tau\left(\nabla f(x^k) + K^T y^k\right) \\ y^{k+1} &= \mathbf{prox}_{\sigma h^*}\left(y^k + \sigma K \tilde{x}^{k+1}\right) \\ x^{k+1} &= \tilde{x}^{k+1} - \tau K^T (y^{k+1} - y^k), \end{aligned} \tag{Algorithm LV}$$

and its relaxed variant

$$\begin{aligned} y^{k+1} &= (1-\rho_k) y^k + \rho_k \, \mathbf{prox}_{\sigma h^*}(y^k + \sigma K \tilde{x}^{k+1}) \\ x^{k+1} &= (1-\rho_k) x^k + \rho_k (\tilde{x}^{k+1} - \tau K^T (y^{k+1} - y^k)), \end{aligned}$$

which converges for (σ, τ) such that

$$\frac{1}{\tau} > \frac{L_f}{2} \quad \text{and} \quad \frac{1}{\tau\sigma} > \|K\|_2^2. \tag{11}$$

Its inertial variant is considered in Section 5. The second class is due to Condat (2013):

$$\begin{aligned} x^{k+1} &= x^k - \tau\left(\nabla f(x^k) + K^T y^k\right) \\ \tilde{x}^{k+1} &= 2x^{k+1} - x^k \\ y^{k+1} &= \mathbf{prox}_{\sigma h^*}(y^k + \sigma K \tilde{x}^{k+1}), \end{aligned} \tag{Algorithm CO}$$

and its relaxed variant, as that for Algorithm LV, which converges for

$$\frac{1}{\tau} > \frac{L_f}{2} \quad \text{and} \quad \left(\frac{1}{\tau} - \frac{L_f}{2}\right)\frac{1}{\sigma} > \|K\|_2^2. \tag{12}$$

Table 1: Choice of f, K, h for various types of losses and penalties for (3). $P_\lambda^{\ell_s}(v)$ is the projection of v onto ℓ_s ball of radius λ.

	strongly smooth loss	not strongly smooth loss		generalized lasso penalty	overlapping ℓ_1-ℓ_q group lasso penalty
$f(x)$	$\sum_{i=1}^n l_i(a_i^T x; b_i)$	0	$H(z)$	$\lambda\|z\|_1$	$\sum_{i=1}^{\mathcal{G}} \lambda_i \|z_{g_i}\|_q$
K	D	$[D^T, A^T]^T$			(g_i: group index, $i = 1, \ldots, \mathcal{G}$)
$h^*(y)$	$H^*(y)$	$H^*(\bar{y}) + \sum_{i=1}^n l_i^*(\tilde{y}_i; b_i)^\dagger$	$\mathbf{prox}_{\sigma H^*}$	$P_\lambda^{\ell_\infty}(z)$	$\left[P_{\lambda_1}^{\ell_s}(z_{g_1})^T, \ldots, P_{\lambda_{\mathcal{G}}}^{\ell_s}(z_{g_{\mathcal{G}}})^T\right]^{T\ddagger}$
$\mathbf{prox}_{\sigma h^*}(z)$	$\mathbf{prox}_{\sigma H^*}(z)$	$(\mathbf{prox}_{\sigma H^*}(\bar{z}),$ $\mathbf{prox}_{\sigma l^*(\cdot,b_i)}(\tilde{z}_i)_{i=1}^n)^\dagger$	D	sparse	membership

$^\dagger y = (\bar{y}, \tilde{y}), z = (\bar{z}, \tilde{z})$ $\quad ^\ddagger 1/q + 1/s = 1$

This condition is slightly more restrictive than (11); an inertial variant is also considered in Section 5. Algorithm CO can be considered as a variant of the primal-dual hybrid gradient method (Zhu and Chan, 2008; Esser et al., 2010) where the evaluation of $\mathbf{prox}_{\tau f}$, which can be costly, is replaced by that of $I - \tau \nabla f$; see Appendix A for details.

Table 1 illustrates specific choices of f, K, and h^* for typical losses and penalties. Note non-smooth losses, e.g., hinge loss, are allowed; see Appendix B.

4 Unification

In this section we provide a unified treatment to Algorithms LV and CO from the perspective of monotone operator theory. It can be shown that both Algorithms LV and CO are instances of preconditioned forward-backward splitting (7). In particular, the preconditioners for Algorithms LV and CO are given by

$$M_{\mathsf{LV}} = \begin{bmatrix} \frac{1}{\tau} I & \\ & \frac{1}{\sigma} I - \tau K K^T \end{bmatrix} \text{ and } M_{\mathsf{CO}} = \begin{bmatrix} \frac{1}{\tau} I & -K^T \\ -K & \frac{1}{\sigma} I \end{bmatrix},$$

respectively. The operator to split is $T = F + G$, where

$$F = \begin{bmatrix} 0 & K^T \\ -K & \partial h^* \end{bmatrix} \quad \text{and} \quad G = \begin{bmatrix} \nabla f & 0 \\ 0 & 0 \end{bmatrix}. \tag{13}$$

Note that using the operator T, the optimality condition (9) for (1) can be written as an inclusion problem

$$\begin{bmatrix} 0 \\ 0 \end{bmatrix} \in T(z^\star) = \begin{bmatrix} \nabla f & K^T \\ -K & \partial h^* \end{bmatrix} \begin{bmatrix} x^\star \\ y^\star \end{bmatrix}. \tag{14}$$

Thus the minimization problem (1) can be understood as finding $z^\star = (x^\star, y^\star) \in \mathbb{R}^p \times \mathbb{R}^l$ satisfying (14). From Section 2, we see that T is maximally monotone.

Combettes et al. (2014) point out that Algorithm LV can be interpreted as a forward-backward splitting. Algorithm CO is analyzed from this perspective by Condat (2013).

In Combettes et al. (2014) and Davis (2015), these algorithms are treated as two separate classes of preconditioning. We provide a simple derivation in Appendix C.

This operator splitting point of view greatly simplifies the convergence analysis of Algorithm LV, and also promotes a unified understanding of both algorithm classes. First we recover the condition for convergence (11) by Chen et al. (2013). Operator F is maximally monotone because it is the sum of the subdifferential of a closed proper convex function $\tilde{h}(x, y) = h^*(y)$ and a skew-symmetric matrix. The following proposition provides a cocoercivity constant for $M^{-1}G$, larger than mentioned in Section 2.

Proposition 2. *Let M be a symmetric, positive definite matrix in $\mathbb{R}^{(p+l)\times(p+l)}$ and G as given in* (13)*. Then, for $\mu > 0$ such that*

$$\|(x, 0)\|^2_{M^{-1}} \leq (1/\mu)\|x\|_2^2, \quad \forall x \in \mathbb{R}^p, \tag{15}$$

operator $M^{-1}G$ is μ/L_f-cocoercive in $\langle\cdot,\cdot\rangle_M$.

The following fact is also useful for interpreting the forthcoming convergence conditions with respect to (2), as well as the analysis of convergence rates in the sequel.

Proposition 3. *For a symmetric, positive definite matrix M and a given $L_f > 0$, the following are equivalent.*

1. *For all $x \in \mathbb{R}^p$, there exists $\mu > L_f/2$ such that $\|(x, 0)\|^2_{M^{-1}} \leq (1/\mu)\|x\|_2^2$.*
2. *There holds*

$$M \succ \begin{bmatrix} \frac{L_f}{2} I & \\ & 0 \end{bmatrix}. \tag{16}$$

3. *There exist $\nu > L_f/2$ and $\epsilon > 0$ such that*

$$M \succeq \begin{bmatrix} \nu I & \\ & \epsilon I \end{bmatrix}.$$

It is natural to set μ^{-1} in condition (15) as the largest eigenvalue of the (1,1) block of the 2×2 block matrix M^{-1}. The class of algorithms of interest has $t = 1$ for (8), and $M^{-1}G$ is μ/L_f-cocoercive in $\langle\cdot,\cdot\rangle_M$ for such μ. Then, $t \in (0, 2\gamma)$ in Proposition 1 is equivalent to requiring $\mu > L_f/2$, thus condition (16). For M_{LV}, μ can be chosen as $1/\tau$. and Algorithm LV converges if $\frac{1}{\tau} > \frac{L_f}{2}$ and $\frac{1}{\tau\sigma} > \|K\|_2^2$, recovering condition (11). The latter condition comes from $M_{\mathsf{LV}} \succ 0$. Furthermore, we see that for $\{\rho_k\} \subset [0, \delta]$ satistying $\sum_{k=0}^{\infty} \rho_k(\delta - \rho_k)$, where $\delta = 2 - (\tau L_f)/2$, the relaxed variant (8) with $M = M_{\mathsf{CO}}$ also converges.

4.1 Relation between Algorithms LV and CO

As claimed, there is an intimate relation between the two class of algorithms. We utilize this to show that Algorithm CO can be derived from Algorithm LV, and recover the condition for convergence (12).

The "LDL" decomposition of M_{CO} reveals that

$$\begin{aligned} M_{\mathsf{CO}} &= \begin{bmatrix} I & \\ -\tau K & I \end{bmatrix} \begin{bmatrix} \frac{1}{\tau} I & \\ & \frac{1}{\sigma} I - \tau K K^T \end{bmatrix} \begin{bmatrix} I & -\tau K^T \\ & I \end{bmatrix} \\ &= L M_{\mathsf{LV}} L^T. \end{aligned} \tag{17}$$

Then it is easy to see that Algorithm CO

$$z^{k+1} = (I + M_{\mathsf{CO}}^{-1} F)^{-1} (I - M_{\mathsf{CO}}^{-1} G)(z^k)$$

is equivalent to

$$L^T z^{k+1} = (I + M_{\mathsf{LV}}^{-1} \tilde{F})^{-1} (I - M_{\mathsf{LV}}^{-1} \tilde{G})(L^T z^k), \tag{18}$$

where $\tilde{F} = L^{-1} F L^{-T}$ and $\tilde{G} = L^{-1} G L^{-T}$. Letting $w = L^T z$, *we see that Algorithm CO is in fact Algorithm LV applied to the linearly transformed variable* w *by splitting the similarly transformed operator* $L^{-1} T L^{-T}$ *by* $\tilde{F}$ *and* $\tilde{G}$. Maximal monotonicity of $M_{\mathsf{LV}}^{-1} \tilde{F}$ is guaranteed (Bauschke and Combettes, 2011, Theorem 24.5), and cocoercivity of G is preserved in $M_{\mathsf{LV}}^{-1} \tilde{G}$ as well:

Proposition 4. *$M_{\mathsf{LV}}^{-1} \tilde{G}$ is $(1/\tau - \sigma \|K\|_2^2)/L_f$-cocoercive with respect to $\|\cdot\|_{M_{\mathsf{LV}}}$.*

Thus Proposition 1 holds with $t = 1$ and $\gamma = (1/\tau - \sigma\|K\|_2^2)/L_f$, and Algorithm CO converges if $\frac{1}{\tau} > \frac{L_f}{2}$ and $\left(\frac{1}{\tau} - \frac{L_f}{2}\right) \frac{1}{\sigma} > \|K\|_2^2$, i.e., (12). With respect to the untransformed sequence $\{z^k\}$, observe that $M_{\mathsf{CO}}^{-1} G$ is also $(1/\tau - \sigma\|K\|_2^2)/L_f$-cocoercive (with respect to $\|\cdot\|_{M_{\mathsf{CO}}}$). In light of (18), it is natural to measure convergence using the metric $\|L^T \cdot\|_{M_{\mathsf{LV}}}$, and this metric coincides with $\|\cdot\|_{M_{\mathsf{CO}}}$. For the relaxed variant, we obtain $\delta = 2 - \frac{L_f}{2} \cdot \frac{1}{1/\tau - \sigma\|K\|_2^2}$ for the range of relaxation coefficients $\{\rho_k\} \subset [0, \delta]$ that guarantees convergence.

4.2 Unified algorithm class

The relation between the two algorithms suggests a more general family of preconditioners, namely

$$M = \tilde{L} M_{\mathsf{LV}} \tilde{L}^T = \begin{bmatrix} \frac{1}{\tau} I & C^T \\ C & \frac{1}{\sigma} I + \tau (C C^T - K K^T) \end{bmatrix}, \tag{19}$$

where $\tilde{L}$ replaces $(2,1)$ block of L in (17) by τC. In particular, if $CK^T = KC^T$, then (7) yields the following iteration:

$$\begin{aligned} y^{k+1} &= \mathbf{prox}_{\sigma h^*}(\sigma K x^k + \sigma\tau(C-K)\nabla f(x^k) \\ &\quad + (I + \sigma\tau K(C-K)^T) y^k) \\ x^{k+1} &= x^k - \tau(\nabla f(x^k) - C^T y^k + (C+K)^T y^{k+1}). \end{aligned} \tag{20}$$

Condition $CK^T = KC^T$ is satisfied if and only if $C = US\Sigma^{-1}V^T + N\bar{V}^T$, where U, V, and Σ are from the reduced singular value decomposition of $K = U\Sigma V^T$ so that Σ is an $r \times r$ positive diagonal matrix where $r = \mathrm{rank}(K)$; $\bar{V}$ is such that $\tilde{V} = [V, \bar{V}]$ is orthogonal; S is symmetric, and N is arbitrary. A simple choice is $S = \kappa\Sigma^2$ for some $\kappa \in \mathbb{R}$ and $N = 0$, yielding $C = \kappa K$. Choosing $\kappa = 0$ and -1 respectively recovers Algorithms LV and CO; for $\kappa = 1$, we have

$$\begin{aligned} y^{k+1} &= \mathbf{prox}_{\sigma h^*}(\sigma K x^k + y^k) \\ x^{k+1} &= x^k - \tau\nabla f(x^k) - \tau K^T(2y^{k+1} - y^k), \end{aligned}$$

which is a dual version of Algorithm CO (Condat, 2013, Algorithm 3.2). Another choice is to set $S = \pm\Sigma^2$ and N so that NN^T is diagonal. In this case $CC^T - KK^T$ reduces to a diagonal matrix, $C = [\bar{K}, N]\tilde{V}$ where $\bar{K}$ is the first r columns of $K\tilde{V}$. If the eigenspace of K^TK is well-known and multiplication with $\bar{V}$ can be computed fast, e.g. the discrete cosine tranform (DCT) matrix for the fused lasso on a regular grid, this choice may be useful.

The same reasoning leading to (12) yields the range of parameters for convergence $1/\tau > L_f/2$ and

$$\left(\frac{1}{\tau} - \frac{L_f}{2}\right)\frac{1}{\sigma} - \frac{\tau L_f}{2}\|C\|_2^2 > \left(1 - \frac{\tau L_f}{2}\right)\|K\|_2^2. \tag{21}$$

Especially when $C = \kappa K$, $\kappa \in [-1, 1]$ can be used to trade-off the range of step size parameters for which (20) converges and an additional vector multiplication with K^T. The range for convergence shrinks gradually from $|\kappa| = 0$ (LV) to 1 (CO); see Figure 1 for visualization.

Remark 1. *Condat (2013) also considers an extenstion of* (1)*, which minimizes the three-function sum $f(x) + g(x) + h(Kx)$, with g closed, proper, and convex. In this case, the first line of Algorithm CO is replaced by $\mathbf{prox}_{\tau g}(x^k - \tau(\nabla f(x^k) + K^T y^k))$. Let us call this extension Algorithm CO+. This algorithm is still preconditioned forward-backward splitting with preconditioner M_{CO}, where the zero in the (1,1) block of operator F is replaced by ∂g, and converges under* (12)*. For this extended F,* (20) *generates a feasible algorithm only when $C = \pm K$, i.e., Algorithm CO+ and its dual. Nevertheless, Algorithm LV also has a three-function extension (Chen et al., 2016).*

Remark 2. *For related preconditioners in which $\frac{1}{\tau}I$, $\frac{1}{\sigma}I$ in M_{LV} and M_{CO} are replaced by Σ_1, $\Sigma_2 \succ 0$ (Pock and Chambolle, 2011; Combettes et al., 2014), the same relation as* (17) *holds. Consequently, a family of preconditioners like* (19) *can be constructed to generate feasible forward-backward splitting algorithms.*

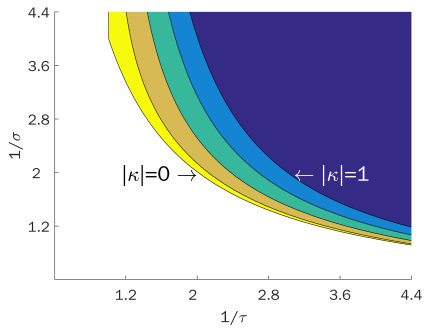

Figure 1: Range of convergence for $(1/\sigma, 1/\tau)$. Boundaries correspond to $|\kappa| = 0$, 0.25, 0.5, 0.75, 1.

5 Rates of convergence

In this section we analyze the rate of convergence of the preconditioned forward-backward splitting algorithm (8) in which F and G are given in (13); M is a symmetric, positive definite matrix in $\mathbb{R}^{(p+l)\times(p+l)}$. As shown in the previous section, an appropriate choice of M yields a feasible algorithm, e.g. Algorithm LV, Algorithm CO, and (20).

We also consider an inertial variant of (8):

$$\tilde{z}^k = z^k + \beta_k(z^k - z^{k-1}) \tag{22a}$$

$$z^{k+1} = R_{M^{-1}F}(I - M^{-1}G)\tilde{z}^k, \tag{22b}$$

where $\{\beta_k\}$ is an external sequence; $\beta_k \equiv 0$ reduces (22) to (8) with $\rho_k \equiv 1$. This variant is motivated by Nesterov's acceleration technique (Nesterov, 1983); see Lorenz and Pock (2015).

Duality gap $G(x, y) = \mathcal{L}(x, y^\star) - \mathcal{L}(x^\star, y)$ can be used to measure the convergence of the objective value, because $G(x^k, y^k) = 0$ guarantees that the pair (x^k, y^k) is a primal-dual solution to (2).

The following lemma sets a basis for establishing rates of convergence of iteration (8) and (22).

Lemma 1. *For operators F, G in* (13)*, a positive definite matrix $M \in \mathbb{R}^{(p+l)\times(p+l)}$, and $\rho \in (0, 2)$, consider a relation $z^+ = (I + M^{-1}F)^{-1}(I - M^{-1}G)z^-$, $z_\rho = (1-\rho)z^- + \rho z^+$. Write $z_\rho = (x_\rho, y_\rho)$, $z^+ = (x^+, y^+)$, $z^- = (x^-, y^-)$, all in*

$\mathbb{R}^p \times \mathbb{R}^l$. *Then the following holds:*

$$\begin{aligned} &2\rho(\mathcal{L}(x^+, y) - \mathcal{L}(x, y^+)) \\ &\leq \|z^- - z\|_M^2 - \|z_\rho - z\|_M^2 + (1 - 2/\rho)\|z^- - z_\rho\|_M^2 \\ &\quad + (L_f/\rho)\|x^- - x_\rho\|_2^2, \ \ \forall z = (x, y) \in \mathbb{R}^p \times \mathbb{R}^l. \end{aligned}$$

5.1 Preconditioned forward-backward splitting

Ergodic convergence. First we establish an $O(1/N)$ ergodic convergence rate of the duality gap evaluated for an average of the first N terms of the sequence $\{\tilde{z}_k\}$:

Theorem 1. *Let M be the symmetric, positive definite matrix in iteration* (8) *and μ satisfy* (15). *Let $\alpha = (2\mu)/(4\mu - L_f)$ and write $z^k = (x^k, y^k)$, $\tilde{z}^k = (\tilde{x}^k, \tilde{y}^k)$. Let $\bar{x}^N = \sum_{k=0}^N \rho_k \tilde{x}^k / \sum_{k=0}^N \rho_k$ and $\bar{y}^N = \sum_{k=0}^N \rho_k \tilde{y}^k / \sum_{k=0}^N \rho_k$. Also let $\bar{\rho} = \sup_{k \geq 0} \rho_k > 0$. If $\mu > L_f/2$ and $\{\rho_k\}$ is chosen so that $0 < \rho_k < 1/\alpha$ for all k, then the following holds for all $z = (x, y) \in \mathbb{R}^p \times \mathbb{R}^l$:*

$$\begin{aligned} &2\left(\mathcal{L}(\bar{x}^N, y) - \mathcal{L}(x, \bar{y}^N)\right) \\ &\qquad \leq \frac{1}{\sum_{k=0}^N \rho_k} \left(\|z^0 - z\|_M^2 + \frac{\alpha L_f}{(1 - \alpha\bar{\rho})\lambda_{\min}(M)} \|z^0 - z^\star\|_M^2 \right), \end{aligned}$$

where $z^\star$ is a point in $\mathbf{zer}\,(F + G)$.

Choosing $(x, y) = (x^\star, y^\star) = z^\star$, we have $O(1/N)$ ergodic convergence rates for the duality gap.

Now let $\mathcal{F}(x) = f(x) + h(Kx)$ be the primal objective function and $\mathcal{F}^\star$ be the primal optimal value. For an important class of the penalty functions h including those for the generalized and group lasso, the following rate for primal suboptimality holds.

Corollary 1. *Assume the conditions for Theorem 1. If $\mathbf{dom}(h) = \mathbb{R}^l$, i.e., $h(\cdot)$ does not take the value $+\infty$, then there exists a constant C_1 independent of N such that*

$$0 \leq \mathcal{F}(\bar{x}^N) - \mathcal{F}^\star \leq C_1 / (\sum_{k=1}^N \rho_k) \ \text{ for all } N.$$

Non-ergodic convergence. Sometimes a non-ergodic solution sequence $\{z_k\}$ or $\{\tilde{z}_k\}$ is preferred as it tends to preserve the desired structural properties better than the ergodic counterpart. Thus it is worth investigating the worst-case performance of the non-ergodic sequence.

Theorem 2. *Let M be the symmetric, positive definite matrix in iteration* (8) *and satisify*

$$M \succeq \begin{bmatrix} \nu I & \\ & \epsilon I \end{bmatrix} \tag{23}$$

for some $\nu > L_f/2$ and $\epsilon > 0$. Let $\alpha = 2\nu/(4\nu - L_f)$ and write $z^k = (x^k, y^k)$, $\tilde{z}^k = (\tilde{x}^k, \tilde{y}^k)$. If $\{\rho_k\}$ is chosen so that $0 < \rho_k < 1/\alpha$ for all k and $\tau = \inf_{k\geq 0} \rho_k(1 - \alpha\rho_k) > 0$, then the following holds for all $z = (x, y) \in \mathbb{R}^p \times \mathbb{R}^l$.

$$\begin{aligned} &\mathcal{L}(\tilde{x}^k, y) - \mathcal{L}(x, \tilde{y}^k) \\ &\quad \leq \|z^0 - z^\star\|_M(\|z^0 - z^\star\|_M + \|z^\star - z\|_M)/(\sqrt{\tau(k+1)}) \end{aligned}$$

and $\mathcal{L}(\tilde{x}^k, y) - \mathcal{L}(x, \tilde{y}^k) = o(1/\sqrt{k+1})$.

With an essentially same proof as that of Corollary 1, we have the following result.

Corollary 2. *Assume the conditions for Theorem 2. If $\mathbf{dom}(h) = \mathbb{R}^l$, then there exists a constant C_2 independent of k such that*

$$0 \leq \mathcal{F}(\tilde{x}^k) - \mathcal{F}^\star \leq C_2/\sqrt{k+1}$$

for all k and $\mathcal{F}(\tilde{x}^k) - \mathcal{F}^\star$ is $o(1/\sqrt{k+1})$.

By Proposition 3, the conditions for Theorems 1 and 2 are both equivalent to (16). Hence the rates of convergence results in this section hold for *any* M satisfying (16). This condition coincides with that of Lorenz and Pock (2015) for (7) to converge, with no rate given, when F and G are given in (13). In particular, for M in (19), this condition reduces to (21). *Thus, for the entire range of (σ, τ) for which* (7)*or* (8) *to converge, we have established an $O(1/N)$ ergodic and an $o(1/\sqrt{k+1})$ non-ergodic convergence rates for objective values.* For Algorithm LV ($M = M_{\mathsf{LV}}$), Loris and Verhoeven (2011) obtain an $O(1/N)$ ergodic convergence rate for $f(x) = \frac{1}{2}\|Ax - b\|_2^2$. For general f, Chen et al. (2013) show that Algorithm LV and its relaxed variant converge under (11), but the rate is given only for strongly convex f and full row rank K. This special case is not very interesting in statistical learning applications in which f is almost always not strongly convex. Our result for the rates of convergence for Algorithm LV and its variants (including the inertial one in the next subsection) without this impractical assumption appears to be novel. For Algorithm CO ($M = M_{\mathsf{CO}}$), our result for the rates is an improvement by a factor of 2 from the analysis by Chambolle and Pock (2016, Theorems 1 and 2). Thus we close the gap between the condition for convergence and that for the rate.

Remark 3. *An inspection of the proof of Lemma 1 asserts that the results of this section also holds for the extended F (see Remark 1). Thus we close the gap for Algorithm CO+, the three-function extenstion of Algorithm CO, as well.*

Remark 4. *Davis (2015, Proposition 5.3) analyzes both ergodic and non-ergodic rates for general F and G, under the condition $M \succeq \underline{\lambda} I$ for some $\underline{\lambda} > 0$. When applied to* (8)*, a range smaller than that is allowed by* (16) *is obtained. This is mainly because of the special structure of G as exploited in Propositions 2 and 3.*

Remark 5. *In Section 7, we demonstrate that the actual performance of the non-ergodic sequence is much better than the theoretical $o(1/\sqrt{k+1})$ rate.*

5.2 Inertial variant

For the inertial variant (22), we present an $O(1/N)$ ergodic rate with minimial assumptions.

Theorem 3. *Let M be the symmetric, positive definite matrix in iteration* (22) *and μ satisfy* (15)*. Write $z^k = (x^k, y^k)$, $\tilde{z}^k = (\tilde{x}^k, \tilde{y}^k)$. Let $\bar{x}^N = (1/N)\sum_{k=1}^N x^k$ and $\bar{y}^N = (1/N)\sum_{k=1}^N y^k$. If $\mu > L_f/2$ and a nondecreasing sequence $\{\beta_k\} \subset [0, \beta]$, $\bar{\beta} \leq 1/3$, is chosen so that there exists an $\eta > 0$ with*

$$(1 - 3\beta_k - \eta)M \succeq \frac{(1-\beta_k)^2}{2} \begin{bmatrix} L_f I & \\ & 0 \end{bmatrix}, \tag{24}$$

for all k, then for the sequence $\{z^k\}$ generated by (22) *with $z^0 = z^{-1} \in \mathbb{R}^p \times \mathbb{R}^l$, the following holds:*

$$\begin{aligned} &\mathcal{L}(\bar{x}^N, y) - \mathcal{L}(x, \bar{y}^N) \\ &\leq \frac{1-\beta_0}{2N}\left(\|z^0 - z\|_M^2 + \frac{2(1+\bar{\beta})^2 L_f}{9(1-\bar{\beta})\eta\lambda_{\min}(M)}\|z^0 - z^\star\|_M^2\right), \end{aligned}$$

where $z^\star$ is a point in $\mathbf{zer}\,(F+G)$.

Again by Proposition 3, the condition for Theorem 3 is equivalent to (23) in Theorem 2 for some $\nu > L_f/2$ and $\epsilon > 0$. Thus condition (24) is satisfied if

$$(1 - 3\beta_k - \eta)\begin{bmatrix} \nu I & \\ & \epsilon I \end{bmatrix} \succeq (1-\beta_k)^2 \begin{bmatrix} (L_f/2) I & \\ & 0 \end{bmatrix}.$$

In other words, the inertial version (22) of (8) (hence those of Algorithms LV and CO) converges at $O(1/N)$ ergodic rate for any $\eta \in (0, (9 - 4L_f/\nu)/(L_f/\nu))$, if $\{\beta_k\}$ is nondecreasing with $0 \leq \beta_k \leq \bar{\beta}$ where

$$\bar{\beta} = \min\left(\frac{1-\eta}{3}, 1 - \frac{3 - \sqrt{9 - 4(L_f/\nu) - 2\eta(L_f/\nu)}}{L_f/\nu}\right).$$

This condition is almost identical to that of Lorenz and Pock (2015), under which (22) converges but the rate is missing, except for a slight loss of the upper bound from $(1+\eta)/3$ to $(1-\eta)/3$.

For the inertial variant of Algorithm CO the same $O(1/N)$ ergodic rate is given (Chambolle and Pock, 2016), for a more restrictive range than above. (With a strong convexity assumption on f and h^*, a linear convergence rate is obtained.) The rate of convergence analysis for the inertial variant of Algorithm LV appears to be the first in the literature.

Remark 6. *There is a more sophisticated variant of Algorithm CO that achieves the optimal rate (Chen et al., 2014). Whether the same optimal rate can be attained for a similar variant of Algorithm LV, or more generally,* (8) *with the family of preconditioners M given in* (19) *would be an interesting outlet for the future research.*

6 Feature-splitting distributed computation

Algorithms LV, CO, and (20) preferably with sparse C, only consist of elementwise operations, matrix-vector multiplications and proximity operators, and they do not require quadratic minimization. Hence they are suitable for distributed processing, provided that $\mathbf{prox}_{\sigma h^*}(\cdot)$ is separable in the partitions. As illustrated in Table 1, this is indeed the case for many statistical learning problems.

We consider the case that the number of features p is so large that, for each sample, the features do not fit into the memory. In other words, the data matrix $A = [A^{[1]}, \dots, A^{[N]}]$, where $A^{[i]} \in \mathbb{R}^{n \times p_i}$, $\sum_{i=1}^N p_i = p$, is stored distributedly in N machines. In this case, it is desirable to also split the vectors $x \in \mathbb{R}^p$ conformally and store distributedly, i.e., $x = [x_{[1]}^T, \dots, x_{[N]}^T]^T$, $x_{[i]} \in \mathbb{R}^{p_i}$. For many instances of (1) including the generalized lasso, $l \gtrsim p$, so it is desirable to partition and store the dual variable $y \in \mathbb{R}^l$ likewise. i.e., $y = [y_{[1]}^T, \dots, y_{[N]}^T]^T$, $y_{[i]} \in \mathbb{R}^{l_i}$, $\sum_{i=1}^N l_i = l$. To compute $K^T y$ and Kx efficiently, it is desirable to also distribute rows and columns of K across the machines, i.e., $K^T = [K_{[1]}^T, \dots, K_{[N]}^T]$ and $K = [K^{[1]}, \dots, K^{[N]}]$, where $K_{[i]} \in \mathbb{R}^{l_i \times p}$, and $K^{[i]} \in \mathbb{R}^{l \times p_i}$. Duplicating K does not incur too much cost, as K is typically sparse. Then, we can carry out computation in a distributed fashion, e.g., using a scatter-gather model, as long as machine i has an access to $A_{[i]}$, $K_{[i]}$, and $K^{[i]}$. In such case, communication cost for multiplication with A or A^T is $O(Nn)$. When the number of nonzero elements in K is $O(p)$, which is the case for examples in Table 1, the communication cost for multiplication with K and K^T is $O(p)$, . This model is especially suitable for multi-GPU platforms. In Appendix D, we provide an example for Algorithm LV.

7 Numerical Experiments

In this section, we illustrate the actual convergence behavior of the studied algorithms. Furthermore, we demonstrate scalability of the distributed algorithm in Section 6 on a multi-GPU platform. We implemented all of the studied algorithms discussed to in the

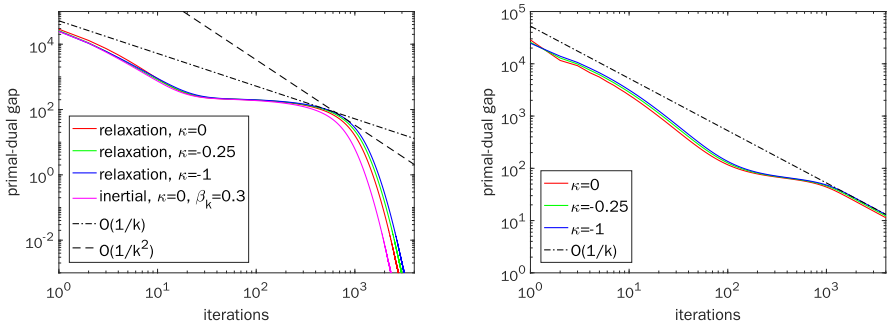

Figure 2: Convergence of relaxation method and inertial method. The left panel shows non-ergodic convergence, and the right panel shows ergodic convergence, measured by the duality gap. $\tau = 0.8(\max)$, $\sigma = 0.8(\max)$.

GPU, using CUDA. The experiment was conducted on the system with two Intel(R) Xeon(R) CPUs (E5-2680 v2, 2.80GHz) with eight Nvidia GTX 980 GPUs with 4 GB of RAM each. We also implemented a distributed version of (20) using CUDA and MPI.

Convergence behavior. We illustrate convergence of iteration (20) with $C = \kappa K$ for overlapping group lasso, based on the setup by Chen et al. (2012). With $n = 5,000$, we defined 1,000 groups of 100 adjacent variables, with 10 overlaps of adjacent groups, thus $p = 90,010$; see Appendix E for the details. The left panel of Figure 2 shows the non-ergodic convergence for various values of κ. We chose ρ_k to be the highest value with theoretical convergence guarantee given by Theorem 2. We can see that the algorithms behave like $O(1/k)$ initially, and then converges faster than $O(1/k^2)$. This is much better than predicted by Theorem 2. Furthermore, they converge faster for smaller values of $|\kappa|$. The right panel illustrates the ergodic convergence of the algorithm. Theorem 1 faithfully describes this kind of convergence. Figure 3 illustrates non-ergodic convergence of the inertial variants. We can see that the smaller $|\kappa|$, the wider the range of β_k that leads the algorithm to convergence. The left panel of Figure 2 also indicates that inertial variants may converge faster than relaxed variants.

Scalability. We computed graph-guided fused lasso with large dimensions following the transcription factor (TF) model of Zhu (2017). This is a simple gene network model with many fully-connected subnetworks. In addition to edges consisting of fully-connected subnetworks, we added random edges between nodes from different subnetworks; see Appendix E for further details. We fixed $n = 5,000$, and varied the dimension p of the data from 120,000 to 1,200,000. Table 2 shows that our distributed algorithm is highly scalable across multiple GPUs. The algorithm runs faster with more GPUs in general, and for the data that do not fit in the memory of small number of GPUs, it can still run with more GPUs.

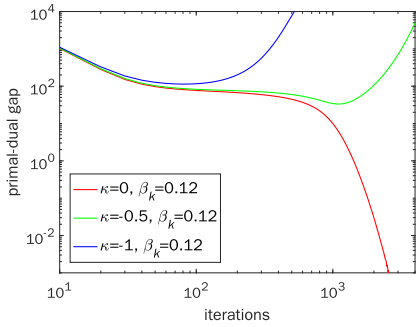

Figure 3: Primal-dual convergence of inertial method for $\beta_k \equiv 0.12$, $\tau = 0.9(\max)$, $\sigma = 0.9(\max)$

Table 2: Scalability of the distributed version of (20) for graph-guided fused lasso model. Time was measured in seconds per 1000 iterations.

	p			
# GPUs	120,000	600,000	960,000	1,200,000
1	83.02	FAILED	FAILED	FAILED
2	55.68	FAILED	FAILED	FAILED
3	46.04	FAILED	FAILED	FAILED
4	44.31	175.04	FAILED	FAILED
5	43.35	164.49	FAILED	FAILED
6	36.66	154.40	FAILED	FAILED
8	37.22	141.18	216.98	270.62

References

Bauschke, H. H. and Combettes, P. L. (2011). *Convex analysis and monotone operator theory in Hilbert spaces*. Springer Science & Business Media.

Beck, A. and Teboulle, M. (2009). A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal on Imaging Sciences*, 2(1):183–202.

Bertsekas, D. P. (2009). *Convex optimization theory*. Athena Scientific.

Boyd, S., Parikh, N., Chu, E., Peleato, B., and Eckstein, J. (2010). Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3(1):1–122.

Boyd, S. P. and Vandenberghe, L. (2004). *Convex optimization*. Cambridge University Press, Cambridge, UK ; New York.

Chambolle, A. and Pock, T. (2011). A first-order primal-dual algorithm for convex problems with applications to imaging. *Journal of Mathematical Imaging and Vision*, 40(1):120–145.

Chambolle, A. and Pock, T. (2016). On the ergodic convergence rates of a first-order primal-dual algorithm. *Mathematical Programming*, 159(1-2):253–287.

Chen, P., Huang, J., and Zhang, X. (2013). A primal-dual fixed point algorithm for convex separable minimization with applications to image restoration. *Inverse Problems*, 29(2):025011.

Chen, P., Huang, J., and Zhang, X. (2016). A primal-dual fixed point algorithm for minimization of the sum of three convex separable functions. *Fixed Point Theory and Applications*, 2016(1).

Chen, X., Lin, Q., Kim, S., Carbonell, J. G., and Xing, E. P. (2012). Smoothing proximal gradient method for general structured sparse regression. *The Annals of Applied Statistics*, 6(2):719–752.

Chen, Y., Lan, G., and Ouyang, Y. (2014). Optimal primal-dual methods for a class of saddle point problems. *SIAM Journal on Optimization*, 24(4):1779–1814.

Combettes, P. L., Condat, L., Pesquet, J.-C., and V, B. C. (2014). A forward-backward view of some primal-dual optimization methods in image recovery. In *2014 IEEE International Conference on Image Processing (ICIP)*, pages 4141–4145. IEEE.

Combettes, P. L. and Vũ, B. C. (2014). Variable metric forwardbackward splitting with applications to monotone inclusions in duality. *Optimization*, 63(9):1289–1318.

Combettes, P. L. and Wajs, V. R. (2005). Signal recovery by proximal forward-backward splitting. *Multiscale Modeling & Simulation*, 4(4):1168–1200.

Condat, L. (2013). A primal-dual splitting method for convex optimization involving lipschitzian, proximable and linear composite terms. *Journal of Optimization Theory and Applications*, 158(2):460–479.

Davis, D. (2015). Convergence rate analysis of primal-dual splitting schemes. *SIAM J. Optim.*, 25(3):1912–1943.

Esser, E., Zhang, X., and Chan, T. F. (2010). A general framework for a class of first order primal-dual algorithms for convex optimization in imaging science. *SIAM Journal on Imaging Sciences*, 3(4):1015–1046.

Friedman, J., Hastie, T., Höfling, H., and Tibshirani, R. (2007). Pathwise coordinate optimization. *The Annals of Applied Statistics*, 1(2):302–332.

Gabay, D. (1983). *Applications of the method of multipliers to variational inequalities*, volume 15, pages 299–331. Elsevier.

Goldstein, T. and Osher, S. (2009). The split Bregman method for L1-regularized problems. *SIAM Journal on Imaging Sciences*, 2(2):323–343.

He, B. and Yuan, X. (2012). Convergence analysis of primal-dual algorithms for a saddle-point problem: From contraction perspective. *SIAM Journal on Imaging Sciences*, 5(1):119–149.

Huber, P. J. (1981). *Robust Statistics*. Wiley, New York, 1 edition.

Liu, J., Yuan, L., and Ye, J. (2010). An efficient algorithm for a class of fused lasso problems. In *Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 323–332. ACM.

Lorenz, D. A. and Pock, T. (2015). An inertial forward-backward algorithm for monotone inclusions. *Journal of Mathematical Imaging and Vision*, 51(2):311–325.

Loris, I. and Verhoeven, C. (2011). On a generalization of the iterative soft-thresholding algorithm for the case of non-separable penalty. *Inverse problems*, 27(12):125007.

Nesterov, Y. (1983). A method of solving a convex programming problem with convergence rate o (1/k2). *Dokl. Akad. Nauk SSSR*, 269(3):543–547.

Nesterov, Y. (2005). Smooth minimization of non-smooth functions. *Mathematical Programming*, 103(1):127–152.

Pock, T. and Chambolle, A. (2011). Diagonal preconditioning for first order primal-dual algorithms in convex optimization. In *Computer Vision (ICCV), 2011 IEEE International Conference on*, pages 1762–1769. IEEE.

Rudin, L. I., Osher, S., and Fatemi, E. (1992). Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena*, 60(1):259–268.

Tibshirani, R., Saunders, M., Rosset, S., Zhu, J., and Knight, K. (2005). Sparsity and smoothness via the fused lasso. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(1):91–108.

Tibshirani, R. J. and Taylor, J. (2011). The solution path of the generalized lasso. *The Annals of Statistics*, 39(3):1335–1371.

Xin, B., Kawahara, Y., Wang, Y., and Gao, W. (2014). Efficient generalized fused lasso and its application to the diagnosis of alzheimer's disease. In *AAAI*, pages 2163–2169.

Ye, G.-B. and Xie, X. (2011). Split bregman method for large scale fused lasso. *Computational Statistics & Data Analysis*, 55(4):1552–1569.

Yuan, M. and Lin, Y. (2006). Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 68(1):49–67.

Zhu, M. and Chan, T. (2008). An efficient primal-dual hybrid gradient algorithm for total variation image restoration. *UCLA CAM Report*, pages 08–34.

Zhu, Y. (2017). An augmented admm algorithm with application to the generalized lasso problem. *Journal of Computational and Graphical Statistics*, (to appear).

A Related Methods

Proximal gradient method. The proximal gradient method is a forward-backward splitting algorithm applied to the primal (1). The first-order optimality condition (9) is given by

$$0 \in \nabla f(x^\star) + K^T \partial h(Kx^\star). \tag{25}$$

Letting $F = K^T \circ \partial h \circ K$ (maximally monotone) and $G = \nabla f$ ($1/L_f$-cocoercive, see Lemma 2 in Appendix F) and applying forward-backward splitting , we obtain $x^{k+1} = \mathbf{prox}_{th\circ K}(x^k - t\nabla f(x^k)) = \arg\min_x f(x^k) + \langle \nabla f(x^k), x - x^k\rangle + \frac{1}{2t}\|x - x^k\|_2^2 + h(Kx)$ which converges for $0 < t < 2/L_f$. When $h(\cdot) = \lambda\|\cdot\|_1$ and $K = I$, this is the celebrated iterative shrinkage-thresholding algorithm (ISTA, Beck and Teboulle, 2009). For this particular h and K being a discrete difference matrix on a graph (fused lasso), evaluating $\mathbf{prox}_{th\circ K}(v)$ is equivalent to solving the fused lasso signal approximator (FLSA, Friedman et al., 2007) that minimizes $\frac{1}{2}\|x - v\|_2^2 + t\lambda\|Kx\|_1$, for which efficient methods have been proposed (Liu et al., 2010; Xin et al., 2014).

However, for general K and other choices of h, e.g., group lasso, evaluation of $\mathbf{prox}_{th\circ K}$ is not trivial. In this case, approximating h by a smooth function can be considered (Nesterov, 2005). Chen et al. (2012) propose to approximate $h(x)$ by its Moreau-Yosida smoothing $\tilde{h}_\mu(x) = \inf_{u+v=x} h(u) + \frac{1}{2\mu}\|v\|_2^2$. In particular, when $h(x) = \lambda\|x\|_1$, $\tilde{h}_\mu(x)$ is the Huber function (Huber, 1981). Because $\tilde{h}_\mu$ is differentiable and $\nabla\tilde{h}_\mu$ is Lipschitz continuous (with constant $1/\mu$), the proximal gradient method is applicable.

Alternating directions method of multipliers (ADMM). ADMM (see, e.g., Boyd et al., 2010) solves (1) using the augmented Lagrangian $\mathcal{L}_t(x, \tilde{x}, y) = f(x) + h(\tilde{x}) + \langle y, Kx - \tilde{x}\rangle + (t/2)\|Kx - \tilde{x}\|_2^2$ in a manner similar to the Gauss-Seidel iteration:

$$\begin{aligned}
x^{k+1} &= \arg\min_x \mathcal{L}_t(x, \tilde{x}^k, y^k) && \text{(26a)}\\
&= \arg\min_x f(x) + (t/2)\|Kx - \tilde{x}^k + (1/t)y^k\|_2^2 \\
\tilde{x}^{k+1} &= \arg\min_{\tilde{x}} \mathcal{L}_t(x^{k+1}, \tilde{x}, y^k) \\
&= \mathbf{prox}_{(1/t)h}(Kx^{k+1} + (1/t)y^k) && \text{(26b)}\\
y^{k+1} &= y^k + t(Kx^{k+1} - \tilde{x}^{k+1}) && \text{(26c)}
\end{aligned}$$

The x-update (26a) is potentially expensive to compute. For example, if $f(x) = \frac{1}{2}\|Ax - b\|_2^2$, then the corresponding update involves solving a linear equation of the form $(A^TA + tK^TK)x = r$. While K is structured and known *a priori*, A is data-dependent in regression problems hence the matrix $A^TA + tK^TK$ is hardly structured. Ye and Xie (2011) propose to solve (26a) iteratively using the preconditioned conjugate gradient method.

ADMM (26) for (1) is equivalent to the Douglas-Rachford operator splitting method applied to the Fenchel dual (Gabay, 1983) :

$$\text{minimize } h^*(y) + f^*(-Ky). \tag{27}$$

The Douglas-Rachford splitting utilizes the equivalence of $0 \in \mathbf{zer}\,(F+G)$ to $z = R_{tG}(\tilde{z})$ with $\tilde{z} \in \mathbf{Fix}\,(2R_{tF}-I)(2R_{tG}-I)$. This relation suggests a KM iteration $\tilde{z}^{k+1} = \frac{1}{2}\tilde{z}^k + \frac{1}{2}(2R_F - I)(2R_G - I)\tilde{z}^k$, $z^{k+1} = R_G(\tilde{z}^{k+1})$. It can be shown that by applying this iteration to (27) with $F = \partial h^*$ and $G = (-K)^T \circ \partial f^* \circ (-K)$, (26) is obtained. Note that evaluation of $R_{tG}(v) = \mathbf{prox}_{tf^*\circ(-K)}(v)$ is not straightforward; this is connected to the difficulty with (26a).

Primal-dual hybrid gradient method (PDHG). PDHG (Zhu and Chan, 2008; Esser et al., 2010) is the following iteration

$$x^{k+1} = \mathbf{prox}_{\tau f}(x^k - \tau K^T y^k) \tag{28a}$$

$$\tilde{x}^{k+1} = 2x^{k+1} - x^k \tag{28b}$$

$$y^{k+1} = \mathbf{prox}_{\sigma h^*}(y^k + \sigma K \tilde{x}^{k+1}), \tag{28c}$$

which converges if $\sigma, \tau > 0$ and $\sigma\tau\|K\|_2^2 < 1$ (Chambolle and Pock, 2011; He and Yuan, 2012). Note that evaluation of $\mathbf{prox}_{th\circ K}$ is replaced by $\mathbf{prox}_{\sigma h^*}$, which is usually simpler; see Table 1 for examples. Nevertheless, similar to ADMM, when f is data-dependent, $\mathbf{prox}_{\tau f}(v)$ is difficult to evaluate: if $f(x) = \frac{1}{2}\|Ax - b\|_2^2$, then (28a) involves solving a linear equation of the form $(tI + A^TA)x = r$.

PDHG can be motivated from several approaches. In particular, (28) can be obtained from ADMM applied to the Fenchel dual (27):

$$\begin{aligned}
\tilde{y}^{k+1} &= \mathbf{prox}_{(1/\tau)f^*}((1/\tau)x^k - K^T y^k) \\
x^{k+1} &= x^k + \tau(K^T y^k + \tilde{y}^{k+1}) \\
y^{k+1} &= \arg\min_y h^*(y) + (\tau/2)\|K^T y + z^{k+1} - (1/\tau)x^{k+1}\|_2^2
\end{aligned}$$

by adding an additional regularization term $\frac{1}{2}\|y - y^k\|^2_{(1/\sigma)I - \tau KK^T}$ to the minimization problem in the y-update (Esser et al., 2010; Chambolle and Pock, 2011; Zhu, 2017).

B Examples of different choices of f, K, h^* in the penalized regression problem (3)

This seciton details what is summarized in Table 1.

Generalized lasso. For the generalized lasso, we can put $h(u) = H(u) = \lambda\|u\|_1$ and thus have $h^*(v) = \delta_{\{\|v\|_\infty \le \lambda\}}(v)$ and $\mathbf{prox}_{\sigma h^*}(v) = P_\lambda^{\ell_\infty}(v) = \min(\max(v, -\lambda), \lambda)$,

where δ_S is the indicator function

$$\delta_S(v) = \begin{cases} 0, & v \in S, \\ +\infty, & \text{otherwise}, \end{cases}$$

and $P_\lambda^{\ell_s}(v)$ is the projection of v onto ℓ_s ball of radius λ; the $\min$ and $\max$ operations are conducted element-wisely. Hence $\mathbf{prox}_{\sigma h^*}$ can be computed in parallel.

Overlapping group lasso. For the ℓ_1-ℓ_q group lasso (Yuan and Lin, 2006) with $\mathcal{G}$ possibly overlapping groups, we can choose $H(z) = \lambda_1\|z_1\|_q + \lambda_2\|z_2\|_q + \cdots + \lambda_\mathcal{G}\|z_\mathcal{G}\|_q$ for $z = (z_1^T, z_2^T, \ldots, z_\mathcal{G}^T)^T$, where $z_i \in \mathbb{R}^p$, $\lambda_i > 0$, $i = 1, 2, \ldots, \mathcal{G}$. Setting $D = [D_1, D_2, \ldots, D_\mathcal{G}]^T$, with D_i being a $p \times p$ diagonal matrix whose diagonal entries are 1 if they belong to group $g_i \subset \{1, 2, \ldots, p\}$, and 0 otherwise, $i = 1, \ldots, G$, we obtain $H(Dx) = \lambda_1\|x_{g_1}\|_q + \lambda_2\|x_{g_2}\|_q + \cdots + \lambda_\mathcal{G}\|x_{g_\mathcal{G}}\|_q$, Then we have $h^*(v) = \delta_{\{\|v_{g_1}\|_s \le \lambda_1\}}(v) + \cdots + \delta_{\{\|v_{g_\mathcal{G}}\|_s \le \lambda_\mathcal{G}\}}(v)$, where s satisfies $1/s + 1/q = 1$ ($s = \infty$ if $q = 1$), and $\mathbf{prox}_{\sigma h^*}(v) = \left[P_{\lambda_1}^{\ell_s}(v_{g_1})^T, \ldots, P_{\lambda_\mathcal{G}}^{\ell_s}(v_{g_\mathcal{G}})^T\right]^T$. Thus $\mathbf{prox}_{\sigma h^*}$ can be computed in group-wise parallel.

Not strongly smooth loss. When the loss function l_i is not strongly smooth but still closed, proper, and convex, we can further dualize this function so that $\sum_{i=1}^n l_i(a_i^T x; b_i) = \sup_{w \in \mathbb{R}^n} \langle Ax, w\rangle - \sum_{i=1}^n l_i^*(w_i; b_i)$, yielding

$$\min_x \sup_{y,w} \langle Dx, y\rangle + \langle Ax, w\rangle - \left(\sum_{i=1}^n l_i^*(w_i; b_i) + H^*(y)\right).$$

In terms of (2), $f(x) \equiv 0$, $K = [D^T, A^T]^T$, $h^*(y, w) = H^*(y) + \sum_{i=1}^n l_i^*(w_i; b_i)$ in (2). Because h^* is separable in y and w, we have $(y^+, w^+) = \mathbf{prox}_{\sigma h^*}(\overline{y}, \overline{w})$ such that $y^+ = \mathbf{prox}_{\sigma H^*}(\overline{y})$ and $w_i^+ = \mathbf{prox}_{\sigma l_i^*(\cdot; b_i)}(\overline{w}_i)$, $i = 1, \ldots, n$. For example, in linear support vector machine, the proximity operator for the hinge loss $l_i(\cdot, b_i) = \max(0, 1 - b_i\cdot)$ is given by $\mathbf{prox}_{\sigma l_i^*}(w_i) = \max(\min(w_i - \sigma b_i, 0), -b_i)$. Thus computation of $\mathbf{prox}_{\sigma h^*}$ can be conducted in parallel for elements of w. Note this dualization technique can be also applied to PDHG, which in this case coincides with Algorithm CO. For smooth losses, e.g., quadratic loss, however, both Algorithms LV and CO can proceed without dualization, which increases the dimension of the dual variable by n; PDHG may have to dualize the loss terms unless the proximity operator of $f(x) = \sum_{i=1}^n l_i(a_i^T x; b_i)$ is simple to evaluate.

C Derivation of preconditioned forward-backward splitting algorithms

Algorithm LV

Write Algorithm LV as

$$x^{k+1} = x^k - \tau\left(\nabla f(x^k) + K^T y^{k+1}\right) \tag{29}$$

$$\begin{aligned}(I + \sigma\partial h^*)(y^{k+1}) &\ni y^k + \sigma K\tilde{x}^{k+1} \\ &= y^k + \sigma K(I - \tau\nabla f)(x^k) - \sigma\tau KK^T y^k \\ &= \sigma K(I - \tau\nabla f)(x^k) + (I - \sigma\tau KK^T)y^k \end{aligned} \tag{30}$$

Right-multiply σK to both sides of (29) and subtract from (30) to obtain

$$-\sigma K x^{k+1} + (I - \sigma\tau KK^T + \sigma\partial h^*)(y^{k+1}) \ni (I - \sigma\tau KK^T)y^k \tag{31}$$

Combine (29) and (31) to write

$$\begin{bmatrix} \frac{1}{\tau}I & K^T \\ -K & \frac{1}{\sigma}I - \tau KK^T + \partial h^* \end{bmatrix} \begin{bmatrix} x^{k+1} \\ y^{k+1} \end{bmatrix} \ni \begin{bmatrix} \frac{1}{\tau}I - \nabla f & \\ & \frac{1}{\sigma}I - \tau KK^T \end{bmatrix} \begin{bmatrix} x^k \\ y^k \end{bmatrix} \tag{32}$$

Let

$$M_{\mathsf{LV}} = \begin{bmatrix} \frac{1}{\tau}I & \\ & \frac{1}{\sigma}I - \tau KK^T \end{bmatrix}, \quad F = \begin{bmatrix} 0 & K^T \\ -K & \partial h^* \end{bmatrix}, \quad G = \begin{bmatrix} \nabla f & 0 \\ 0 & 0 \end{bmatrix}.$$

Then we see that (32) is

$$(M_{\mathsf{LV}} + F)(z^{k+1}) \ni (M_{\mathsf{LV}} - G)(z^k),$$

where $z^k = (x^k, y^k)$, which in turn is equivalent to

$$z^{k+1} = (I + M_{\mathsf{LV}}^{-1}F)^{-1}(I - M_{\mathsf{LV}}^{-1}G)(z^k). \tag{33}$$

Note that $F + G = T$, where T is given in (14), the optimality condition for (3).

Algorithm CO

Write Algorithm CO as

$$x^{k+1} = x^k - \tau\left(\nabla f(x^k) + K^T y^k\right) \tag{34}$$

$$\begin{aligned}(I + \sigma\partial h^*)(y^{k+1}) &\ni y^k + \sigma K\tilde{x}^{k+1} \\ &= y^k + \sigma K(2x^{k+1} - x^k) \\ &= y^k - \sigma K x^k + 2\sigma K x^{k+1}, \end{aligned} \tag{35}$$

Inclusion (35) is equivalent to

$$-2\sigma K x^{k+1} + (I + \sigma h^*)(y^{k+1}) \ni -\sigma K x^k + y^k \tag{36}$$

Combine (34) and (36) to write

$$\begin{bmatrix} \frac{1}{\tau} I & 0 \\ -2K & \frac{1}{\sigma} I + \partial h^* \end{bmatrix} \begin{bmatrix} x^{k+1} \\ y^{k+1} \end{bmatrix} \ni \begin{bmatrix} \frac{1}{\tau} I - \nabla f & -K^T \\ -K & \frac{1}{\sigma} I \end{bmatrix} \begin{bmatrix} x^k \\ y^k \end{bmatrix} \tag{37}$$

Let

$$M_{\mathsf{CO}} = \begin{bmatrix} \frac{1}{\tau} I & -K^T \\ -K & \frac{1}{\sigma} I \end{bmatrix}, \quad F = \begin{bmatrix} 0 & K^T \\ -K & \partial h^* \end{bmatrix}, \quad G = \begin{bmatrix} \nabla f & 0 \\ 0 & 0 \end{bmatrix}. \tag{38}$$

Then we see that (37) is

$$(M_{\mathsf{CO}} + F)(z^{k+1}) \ni (M_{\mathsf{CO}} - G)(z^k),$$

where $z^k = (x^k, y^k)$, which in turn is equivalent to

$$z^{k+1} = (I + M_{\mathsf{CO}}^{-1} F)^{-1} (I - M_{\mathsf{CO}}^{-1} G)(z^k).$$

Note that the operators F and G are the same as Algorithm LV.

D Example of the Feature-Splitting Distributed Algorithm

In this section, we illustrate how to construct a distributed algorithm from Algorithm LV. We split p features into N partitions, such that $\sum_{i=1}^N p_i = p$, $A = [A^{[1]}, \ldots, A^{[N]}]$, with $A^{[i]} \in \mathbb{R}^{n \times p_i}$, $K = [K^{[1]}, \ldots, K^{[N]}]$, with $K^{[i]} \in \mathbb{R}^{l \times p_i}$, and $x = [x_{[1]}, \ldots, x_{[N]}]$, with $x_{[i]} \in \mathbb{R}^{p_i}$. The dimension of y is often $O(p)$ in many penalized regression problems, so it is also desirable to split the dual dimension l such that $\sum_{i=1}^N l_i = l$, $K^T = [K_{[1]}, \ldots, K_{[N]}]$ with $K_{[i]} \in \mathbb{R}^{l_i \times p}$, and $y = [y_{[1]}, \ldots, y_{[N]}]$, with $y_{[i]} \in \mathbb{R}^{l_i}$. We store $A^{[i]}$, $K^{[i]}$, $K_{[i]}$, $x_{[i]}$, and $y_{[i]}$ in machine i, $i = 1, \ldots, N$. Duplicating K does not incur too much cost, as K is assumed sparse.

For $f(x) = \sum_{i=1}^n l_i(a_i^T x, b_i)$, $K = D$, and $h = H$, as in (3), we have $\nabla f(x) = \sum_{i=1}^n a_i l_i'(a_i^T x, b_i) = A^T r$, where we set $r = \nabla l(Ax, b) \in \mathbb{R}^n$ with $l(w, b) := [l_1(w_1, b_1), \ldots, l_n(w_n, b_n)]^T$, which can be understood as the residual. Then the itera-

tion of Algorithm LV can be written as follows:

$$\tilde{x}^{k+1} = x^k - \tau(s^k + q^k) \tag{39a}$$
$$t^{k+1} = D\tilde{x}^{k+1} \tag{39b}$$
$$y^{k+1} = \mathbf{prox}_{\sigma H^*}(y^k + \sigma t^{k+1}) \tag{39c}$$
$$s^{k+1} = D^T y^{k+1} \tag{39d}$$
$$x^{k+1} = \tilde{x}^{k+1} - \tau(s^{k+1} - s^k) \tag{39e}$$
$$u^{k+1} = Ax^{k+1} \tag{39f}$$
$$r^{k+1} = \nabla l(u^{k+1}, b) \tag{39g}$$
$$q^{k+1} = A^T r^{k+1}. \tag{39h}$$

Steps (39a), (39e), and (39g) are elementwise operations, and their distribution is straightforward. For computing Step (39b), each worker machine computes $D^{[i]}\tilde{x}_{[i]}$, and send it to the master machine to compute the sum: $D\tilde{x} = \sum_{i=1}^N D^{[i]}\tilde{x}_{[i]}$. Each $D^{[i]}x_{[i]}$ is in $\mathbb{R}^l$, but as D is sparse, only a few of the elements of $D^{[i]}x_{[i]}$ are nonzero. It is sufficient to send only those elements. Hence, the communication cost in this step is $O(p)$. Distributing $\{t^k\}$ again costs $O(p)$. Step (39c) can easily be computed in distributed fashion, as long as computation of $\mathbf{prox}_{\sigma H^*}$ is separable within each partition. Step (39d) is analogous to (39b). For Step (39f), each worker machine computes $A^{[i]}x_{[i]}$, and send it to the master machine, and sum them up: $Ax = \sum_{i=1}^N A^{[i]}x_{[i]}$. The cost for this communication is $O(Nn)$. The sum vector $\{u^k\}$ has dimension n, and we do not need to distribute it. In Step (39h), as the dimension of r is only n, we can broadcast it to all the machines, incurring cost of $O(Nn)$. Then, we compute $A^{[i]T}r$ partitionwise.

E Data Generation

Overlapping group lasso. We generated an overlapping group lasso dataset, based on Chen et al. (2012). We defined 1,000 groups of 100 adjacent variables, with 10 overlaps of adjacent groups. i.e., $g_j = \{90(j-1)+1, \ldots, 90j+10\}$, thus $p = 90,010$, with $n = 5,000$. We set $x_j = (-1)^j \exp(-(j-1)/100)$ for $j = 1, \ldots, \mathcal{G}$. We sampled each element of A from the standard normal distribution, and added Gaussian noise $\epsilon \sim \mathcal{N}(0,1)$ to Ax to generate $b = Ax + \epsilon$, with $n = 5,000$.

Graph-guided fused lasso. The dataset for graph-guided fused lasso experiment was generated following the transcription factor (TF) model of (Zhu, 2017). This is a simple gene network model with k fully connected subnetworks of size r, where each subgroup has one TF with $r-1$ regulatory target genes. Variables corresponding to TFs are sampled i.i.d. from $\mathcal{N}(0,1)$. Variables for target genes are sampled so that

each target gene and the corresponding TF has a bivariate normal distribution with correlation 0.7, and these variables are conditionally independent given the TF. For j-th subnetwork, we chose

$$x_i = \begin{cases} (-1)^{j+1} \left\lfloor \frac{j+1}{2} \right\rfloor & \text{if } j = 1, \ldots, k_a \\ 0 & \text{otherwise} \end{cases}, \quad i = (j-1)r + 1, \ldots, jr,$$

where k_a is the number of active groups. Response b_i is sampled so that $b_i = Ax + \epsilon_i$, with $\epsilon_i \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, 100^2)$. In addition to edges comprising fully connected subnetorks, we added random edges between active variables and inactive variables. For each active variables, we added edges connecting this variable and $k-1$ distinct inactive variables. For our experiment, we set $r = 12$, $k_a = 20$. We selected $k = 10000$, 50000, 80000, and 100000 to generate dataset with $p = 120000$, 600000, 960000, and 1200000, respectively.

F Proofs

We need the following lemma to prove Proposition 2.

Lemma 2 (Combettes and Wajs (2005))**.** *Let ϕ be a convex closed proper function. If $\nabla\phi(z)$ is Lipschitz continuous with Lipschitz constant L, then $\nabla\phi$ is $1/L$-cocoercive and vice versa.*

Proof of Proposition 2.

$$\begin{aligned} \|M^{-1}Gz - M^{-1}Gz'\|_M^2 = \|Gz - Gz'\|_{M^{-1}}^2 &= \|(\nabla f(x) - \nabla f(x'), 0)\|_{M^{-1}}^2 \\ &\leq (1/\mu)\|\nabla f(x) - \nabla f(x')\|_2^2 \\ &\leq (L_f/\mu)\langle \nabla f(x) - \nabla f(x'), x - x' \rangle \\ &= (L_f/\mu)\langle Gz - Gz', z - z' \rangle \\ &= (L_f/\mu)\langle M^{-1}Gz - M^{-1}Gz', z - z' \rangle_M. \end{aligned}$$

Note that we used Lemma 2 in the third line. □

Proof of Proposition 3. We first show that Condition 2 is equivalent to

$$0 \prec M^{-1} \prec \begin{bmatrix} \frac{2}{L_f} I & \\ & \infty \end{bmatrix}, \tag{40}$$

or $z^T M^{-1} z < \frac{2}{L_f}\|x\|_2^2 + \delta_{\{0\}}(y)$ for all $z = (x, y) \neq 0$, where $\delta_S(u)$ is the indicator function, i.e., $\delta_S(u) = 0$ if $u \in S$ and $\delta_S(u) = \infty$ otherwise. To see this, let $g_1(z) = (1/2) z^T M z$ and $g_2(z) = \frac{1}{2} z^T \begin{bmatrix} \frac{L_f}{2} I & \\ & 0 \end{bmatrix} z = \frac{L_f}{4}\|x\|_2^2$. Then Condition 2 ensures that

$g_1(z) > g_2(z)$ for all $z \neq 0$. Take the convex conjugates of g_1 and g_2. Observe that for $w = (w_1, w_2)$, $g_1^*(w) = \sup_z \langle w, z\rangle - (1/2) z^T M z = (1/2) w^T M^{-1} w$ and

$$g_2^*(w) = \sup_z \langle w, z\rangle - g_2(z) = \sup_x \langle w_1, x\rangle - \tfrac{L_f}{4}\|x\|_2^2 + \sup_y \langle w_2, y\rangle = \begin{cases} \frac{1}{L_f}\|w_1\|_2^2, & \text{if } w_2 = 0, \\ \infty, & \text{otherwise.} \end{cases}$$

Conjugacy asserts that $g_1^*(w) \leq g_2^*(w)$, or equivalently

$$0 \prec M^{-1} \preceq \begin{bmatrix} \frac{2}{L_f} I & \\ & \infty \end{bmatrix}.$$

Now for $w = (w_1, 0)$ ($w_1 \neq 0$), $f_1^*(w) = \langle w, \hat{z}\rangle - (1/2)\hat{z}^T M \hat{z} = (1/2) w_1 \bar{M}_{11} w_1$, where

$$M^{-1} = \begin{bmatrix} \bar{M}_{11} & \bar{M}_{12} \\ \bar{M}_{12}^T & \bar{M}_{22} \end{bmatrix}, \quad \hat{z} = M^{-1} w = \begin{bmatrix} \bar{M}_{11} w_1 \\ \bar{M}_{12}^T w_1 \end{bmatrix} \neq 0,$$

because $\bar{M}_{11} \succ 0$. Then

$$\tfrac{1}{2} w_1^T \bar{M}_{11} w_1 = g_1^*(w) = \langle w, \hat{z}\rangle - g_1(\hat{z}) < \langle w, \hat{z}\rangle - g_2(\hat{z}) \leq \sup_z \langle w, z\rangle - g_2(z) = g_2^*(w) = \tfrac{1}{L_f}\|w_1\|_2^2,$$

or $\bar{M}_{11} \prec \frac{2}{L_f} I$. It follows (40). Because both g_1 and g_2 are convex, closed, and proper, the same logic applies to g_1^* and g_2^*, meaning that the above matrix inequality implies Condition 2, establishing the equivalence.

Now Condition 1 implies $(x, 0)^T M^{-1} (x, 0) < \frac{2}{L_f}\|x\|_2^2$ for all $x \neq 0$ and $z^T M^{-1} z < \infty$, implying (40), thus Condition 2. That Condition 2 implies Condition 1 is straightforward, by choosing $1/\mu \in [\lambda_{\max}(\bar{M}_{11}), 2/L_f)$.

Condition 3 is equivalent to

$$0 \prec M^{-1} \preceq \begin{bmatrix} \frac{1}{\nu} I & \\ & \frac{1}{\epsilon} I \end{bmatrix}, \tag{41}$$

thus $(x, 0)^T M^{-1} (x, 0) \leq \frac{1}{\nu}\|x\|_2^2$ where $\nu > L_f/2$. This implies Condition 1. Finally, note that

$$z^T M^{-1} z = x^T \bar{M}_{11} x + 2 x^T \bar{M}_{12}^T y + y^T \bar{M}_{22} y \leq \lambda_{\max}(\bar{M}_{11})\|x\|_2^2 + 2 x^T \bar{M}_{12}^T y + \lambda_{\max}(\bar{M}_{12})\|y\|_2^2,$$

or

$$M^{-1} \preceq \begin{bmatrix} \lambda_{\max}(\bar{M}_{11}) I & \bar{M}_{12} \\ \bar{M}_{12}^T & \lambda_{\max}(\bar{M}_{22}) I \end{bmatrix}.$$

Both $\lambda_{\max}(\bar{M}_{11})$ and $\lambda_{\max}(\bar{M}_{22})$ are positive because $\bar{M}_{11}, \bar{M}_{22} \succ 0$. Then the second inequality in (41) holds if and only if either $\frac{1}{\nu} = \lambda_{\max}(\bar{M}_{11})$, $\bar{M}_{12} = 0$, $\frac{1}{\epsilon} - \lambda_{\max}(\bar{M}_{22}) \geq 0$ or $\frac{1}{\nu} > \lambda_{\max}(\bar{M}_{11})$, $\bar{M}_{12} = 0$, $\frac{1}{\epsilon} - \lambda_{\max} \geq (\frac{1}{\nu} - \lambda_{\max}(\bar{M}_{11}))^{-1}\|\bar{M}_{12}\|_2^2$ (Boyd and Vandenberghe, 2004, Appendix A). Now because Condition 1 implies $\lambda_{\max}(\bar{M}_{11}) \leq \frac{1}{\mu} < \frac{2}{L_f}$, we can choose ν and ϵ so that $\frac{1}{\mu} \leq \frac{1}{\nu} < \frac{2}{L_f}$ and $\frac{1}{\epsilon} \geq \lambda_{\max}(\bar{M}_{22}) + (\frac{1}{\nu} - \lambda_{\max}(\bar{M}_{11}))^{-1}\|\bar{M}_{12}\|_2^2$. This implies (41) and thus Condition 3. □

Proof of Proposition 4. Note that $\|L^T\cdot\|^2_{M_{\mathsf{LV}}} = \langle M_{\mathsf{LV}}L^T\cdot, L^T\cdot\rangle = \|\cdot\|^2_{M_{\mathsf{CO}}}$ and likewise $\langle L^T\cdot, L^T\cdot\rangle_{M_{\mathsf{LV}}} = \langle\cdot,\cdot\rangle_{M_{\mathsf{CO}}}$. Then,

$$\begin{aligned}
&\langle M_{\mathsf{LV}}^{-1}L^{-1}GL^{-T}w - M_{\mathsf{LV}}^{-1}L^{-1}GL^{-T}w', w - w'\rangle_{M_{\mathsf{LV}}} \\
&\quad= \langle (L^TM_{\mathsf{CO}}^{-1}L)(L^{-1}GL^{-T})(L^Tz) - (L^TM_{\mathsf{CO}}^{-1}L)(L^{-1}GL^{-T})(L^Tz'), L^Tz - L^Tz'\rangle_{M_{\mathsf{LV}}} \\
&\quad= \langle L^T(M_{\mathsf{CO}}^{-1}Gz - M_{\mathsf{CO}}^{-1}Gz'), L^T(z-z')\rangle_{M_{\mathsf{LV}}} \\
&\quad= \langle M_{\mathsf{CO}}^{-1}Gz - M_{\mathsf{CO}}^{-1}Gz', z - z'\rangle_{M_{\mathsf{CO}}} \\
&\quad\geq (\mu/L_f)\|M_{\mathsf{CO}}^{-1}Gz - M_{\mathsf{CO}}^{-1}Gz'\|_{M_{\mathsf{CO}}} \\
&\quad= (\mu/L_f)\|L^T(L^{-T}M_{\mathsf{LV}}^{-1}L^{-1}GL^{-T}(L^Tz) - L^{-T}M_{\mathsf{LV}}^{-1}L^{-1}GL^{-T}(L^Tz')\|^2_{M_{\mathsf{LV}}} \\
&\quad= (\mu/L_f)\|M_{\mathsf{LV}}^{-1}L^{-1}GL^{-T}w - M_{\mathsf{LV}}^{-1}L^{-1}GL^{-T}w'\|^2_{M_{\mathsf{LV}}},
\end{aligned}$$

where the inequality and μ come from Proposition 2, as follows. From (17), we see that

$$\begin{aligned}
M_{\mathsf{CO}}^{-1} = L^{-T}M_{\mathsf{LV}}^{-1}L^{-1} &= \begin{bmatrix} I & \tau K \\ & I \end{bmatrix}\begin{bmatrix} \tau I & \\ & (\frac{1}{\sigma}I - \tau KK^T)^{-1}\end{bmatrix}\begin{bmatrix} I & 0 \\ \tau K & I\end{bmatrix} \\
&= \begin{bmatrix} \tau I + \tau^2 K^T(\frac{1}{\sigma}I - \tau KK^T)^{-1}K & \tau K^T(\frac{1}{\sigma}I - \tau KK^T)^{-1} \\ \tau(\frac{1}{\sigma}I - \tau KK^T)^{-1}K & (\frac{1}{\sigma}I - \tau KK^T)^{-1}\end{bmatrix}
\end{aligned}$$

and can choose $\mu = 1/\tau - \sigma\|K\|_2^2$, because $\lambda_{\max}(\tau I + \tau^2K^T(\frac{1}{\sigma}I - \tau KK^T)^{-1}K) = \tau + \tau^2\|K\|_2^2/(1/\sigma - \tau\|K\|_2^2) = \frac{1}{1/\tau - \sigma\|K\|_2^2}$. Therefore $M_{\mathsf{LV}}^{-1}L^{-1}GL^{-T}$ is $(1/\tau - \sigma\|K\|_2^2)/L_f$-cocoercive with respect to $\|\cdot\|_{M_{\mathsf{LV}}}$. □

Proof of Lemma 1. Observe that

$$\begin{aligned}
\|z^- - z\|_M^2 &= \|z^- - z_\rho + z_\rho - z\|_M^2 \\
&= \|z^- - z_\rho\|_M^2 - 2\langle z^- - z_\rho, z - z_\rho\rangle_M + \|z_\rho - z\|_M^2, \qquad (42)
\end{aligned}$$

and, from (8),

$$\begin{aligned}
&z^+ + M^{-1}Fz^+ \ni z^- - M^{-1}Gz^- \\
&\iff z^+ + M^{-1}\begin{bmatrix} 0 & K^T \\ -K & \partial h^*\end{bmatrix}\begin{bmatrix} x^+ \\ y^+\end{bmatrix} \ni z^- - M^{-1}\begin{bmatrix} \nabla f & \\ & 0\end{bmatrix}\begin{bmatrix} x^- \\ y^-\end{bmatrix} \\
&\iff (1/\rho)(z^- - z_\rho) = z^- - z^+ \in M^{-1}\begin{bmatrix} \nabla f(x^-) + K^Ty^+ \\ -Kx^+ + \partial h^*(y^+)\end{bmatrix} \qquad (43)
\end{aligned}$$

Then,

$$\begin{aligned}
\langle z^- - z_\rho, z - z_\rho\rangle_M = \langle \rho(z^- - z^+), z - z_\rho\rangle_M &= \rho\left\langle \begin{bmatrix} \nabla f(x^-) + K^T y^+ \\ -Kx^+ + \partial h^*(y^+) \end{bmatrix}, \begin{bmatrix} x - x_\rho \\ y - y_\rho \end{bmatrix} \right\rangle \\
&= \rho\langle \nabla f(x^-), x - x_\rho\rangle + \rho\langle K^T y^+, x - x_\rho\rangle \\
&\quad + \rho\langle -Kx^+, y - y_\rho\rangle + \rho\langle \partial h^*(y^+), y - y_\rho\rangle \\
&= \rho\langle \nabla f(x^-), x^- - x_\rho\rangle + \rho\langle \nabla f(x^-), x - x^-\rangle + \rho\langle K^T y^+, x - x_\rho\rangle \\
&\quad + \rho\langle -Kx^+, y - y_\rho\rangle + \rho\langle \partial h^*(y^+), y^+ - y_\rho\rangle + \rho\langle \partial h^*(y^+), y - y^+\rangle \\
&\le \rho\langle \nabla f(x^-), x^- - x_\rho\rangle + \rho(f(x) - f(x^-)) + \rho\langle K^T y^+, x - x_\rho\rangle \\
&\quad + \rho\langle -Kx^+, y - y_\rho\rangle + \rho\langle \partial h^*(y^+), y^+ - y_\rho\rangle + \rho(h^*(y) - h^*(y^+)),
\end{aligned} \tag{44}$$

understanding that "$\partial h^*(\cdot)$" represents a subgradient in the corresponding subdifferential. The first and second equalities follow from (43); the last inequality is due to the definition of subgradient. By plugging the inequality (44) in (42) and rearranging terms, we obtain

$$\begin{aligned}
2\rho(\mathcal{L}(x^+, y) - \mathcal{L}(x, y^+)) - \|z^- - z\|_M^2 + \|z_\rho - z\|_M^2 & \\
\le -\|z^- - z_\rho\|_M^2 + 2\rho\langle \nabla f(x^-), x^- - x_\rho\rangle + 2\rho\langle \partial h^*(y^+), y^+ - y_\rho\rangle & \\
- 2\rho\langle K^T y^+, x_\rho\rangle + 2\rho\langle Kx^+, y_\rho\rangle + 2\rho(f(x^+) - f(x^-)) &
\end{aligned} \tag{45}$$

Now it suffices to show that the right-hand side of (45) is less than or equal to $(1 - 2/\rho)\|z^- - z_\rho\|_M^2 + (L_f/\rho)\|x^- - x_\rho\|_2^2$. To see this,

$$\begin{aligned}
\text{(RHS)} &= -\|z^- - z_\rho\|_M^2 + 2\rho\langle \nabla f(x^-), x^- - x_\rho\rangle + 2\rho\langle \partial h^*(y^+), y^+ - y_\rho\rangle \\
&\quad - 2\rho\langle K^T y^+, x_\rho - x^+\rangle - 2\rho\langle K^T y^+, x^+\rangle + 2\rho\langle Kx^+, y_\rho - y^+\rangle + 2\rho\langle Kx^+, y^+\rangle \\
&\quad + 2\rho(f(x^+) - f(x^-)) \\
&= -\|z^- - z_\rho\|_M^2 + 2\rho(f(x^+) - f(x^-) - \langle \nabla f(x^-), x^+ - x^-\rangle) \\
&\quad + 2\rho\langle \nabla f(x^-) + K^T y^+, x^+ - x_\rho\rangle + 2\rho\langle -Kx^+ + \partial h^*(y^+), y^+ - y_\rho\rangle \\
&= -\|z^- - z_\rho\|_M^2 + 2\rho(f(x^+) - f(x^-) - \langle \nabla f(x^-), x^+ - x^-\rangle) \\
&\quad + 2\rho\langle M(z^- - z^+), z^+ - z_\rho\rangle \\
&= -\|z^- - z_\rho\|_M^2 + 2\rho(f(x^+) - f(x^-) - \langle \nabla f(x^-), x^+ - x^-\rangle) \\
&\quad + 2\langle z^- - z_\rho, z^+ - z_\rho\rangle_M \\
&= -\|z^- - z_\rho\|_M^2 + 2\rho(f(x^+) - f(x^-) - \langle \nabla f(x^-), x^+ - x^-\rangle) \\
&\quad + 2(1 - 1/\rho)\langle z^- - z_\rho, z^- - z_\rho\rangle_M \\
&\le (1 - 2/\rho)\|z^- - z_\rho\|_M^2 + \rho L_f\|x^- - x^+\|_2^2 \\
&= (1 - 2/\rho)\|z^- - z_\rho\|_M^2 + (L_f/\rho)\|x^- - x_\rho\|_2^2
\end{aligned}$$

where the third equality follows from (43); the fourth and fifth equalities are from (8); the first inequality is due to the Lipschitz continuity of ∇f; the final equality is again from (8). □

Proof of Theorem 1. From the convexity-concavity of $\mathcal{L}(x, y)$, we have

$$\begin{aligned}\mathcal{L}(\bar{x}^N, y) - \mathcal{L}(x, \bar{y}^N) &\leq \frac{1}{\sum_{k=0}^N \rho_k} \sum_{k=0}^N \rho_k (\mathcal{L}(\tilde{x}^k, y) - \mathcal{L}(x, \tilde{y}^k)) \\ &\leq \frac{1}{2\sum_{k=1}^N \rho_k} \left(\|z^0 - z\|_M^2 + \sum_{k=0}^N \frac{L_f}{\rho_k} \|x^{k+1} - x^k\|_2^2 \right),\end{aligned}$$

where the second inequality comes from Lemma 1 by putting $z^- = z^k$, $z^+ = \tilde{z}^k$, $\rho = \rho_k$, $z_\rho = z^{k+1}$, and noting that $1 < 1/\alpha < 2$ by the assumption $\mu > L_f/2$. Now by Proposition 2 we see that $R_{M^{-1}F}(I - M^{-1}G)$ is α-averaged with respect to $\|\cdot\|_M$, thus by Proposition 1 (2) we have

$$\begin{aligned}\frac{1-\alpha\bar{\rho}}{\alpha} \sum_{k=0}^\infty \frac{1}{\rho_k} \|x^{k+1} - x^k\|_2^2 &\leq \sum_{k=0}^\infty \frac{1-\alpha\rho_k}{\alpha\rho_k} \|x^{k+1} - x^k\|_2^2 \\ &\leq \sum_{k=0}^\infty \frac{1-\alpha\rho_k}{\alpha\rho_k} \|z^{k+1} - z^k\|_2^2 \leq \frac{1}{\lambda_{\min}(M)} \|z^0 - z^\star\|_M^2.\end{aligned}$$

Therefore

$$\begin{aligned}\mathcal{L}(\bar{x}^N, y) - \mathcal{L}(x, \bar{y}^N) &\leq \frac{1}{2\sum_{k=1}^N \rho_k} \left(\|z^0 - z\|_M^2 + \sum_{k=0}^N \frac{L_f}{\rho_k} \|x^{k+1} - x^k\|_2^2 \right) \\ &\leq \frac{1}{2\sum_{k=0}^N \rho_k} \left(\|z^0 - z\|_M^2 + \frac{\alpha L_f}{(1-\alpha\bar{\rho})\lambda_{\min}(M)} \|z^0 - z^\star\|_M^2 \right).\end{aligned}$$

□

Proof of Corollary 1. The proof closely follows that of Loris and Verhoeven (2011, Theorem 1), given for f being quadratic. Because $z^k = (x^k, y^k) \to (x^\star, y^\star) = z^\star \in \mathbf{Fix} T$ where $T = R_{M^{-1}F}(I - M^{-1}G)$, we see $\tilde{z}^k = Tz^k \to z^\star$ and thus $\bar{z}^N = (\bar{x}^N, \bar{y}^N) = (\sum_{k=1}^N \rho_k \tilde{z}^k)/(\sum_{k=1}^N \rho_k) \to z^\star$. Also because $(x^\star, y^\star)$ is a saddle-point of $\mathcal{L}(x, y)$, we have $\mathcal{F}^\star = \mathcal{F}(x^\star) = \mathcal{L}(x^\star, y^\star) \geq \mathcal{L}(x^\star, y)$ for all $y \in \mathbb{R}^l$. Then

$$0 \leq \mathcal{F}(\bar{x}^N) - \mathcal{F}^\star = \mathcal{F}(\bar{x}^N) - \mathcal{L}(x^\star, y^\star) \leq \mathcal{F}(\bar{x}^N) - \mathcal{F}(x^\star, \bar{y}^N) = \sup_{y \in \mathbb{R}^l} \mathcal{L}(\bar{x}^N, y) - \mathcal{L}(x^\star, \bar{y}^N).$$

The $\sup_{y \in \mathbb{R}^l} \mathcal{L}(\bar{x}^N, y) = f(\bar{x}^N) + \sup_{y \in \mathbb{R}^l} \langle K\bar{x}^N, y \rangle - h^*(y)$ is attained at a $\hat{y}^N \in \partial h(K\bar{x}^N)$ because $h^*(\cdot)$ is closed and convex. As $\bar{x}^N$ converges, $K\bar{x}^N$ is bounded independent of N. Now because h is real-valued, it follows that h is locally Lipschitz in the neighborhood of $K\bar{x}^N$ (see, e.g., Bertsekas, 2009, Proposition 5.4.2). Let the local Lipschitz constant be Q. It also follows that $\partial h(K\bar{x}^N)$ is bounded by Q, i.e.

$\|\hat{y}^N\|_2 \leq Q$. Therefore

$$\begin{aligned}
0 \leq \mathcal{F}(\bar{x}^N) - \mathcal{F}^\star &= \mathcal{F}(\bar{x}^N) - \mathcal{L}(x^\star, y^\star) = \sup_{y \in \mathbb{R}^l} \mathcal{L}(\bar{x}^N, y) - \mathcal{L}(x^\star, \bar{y}^N) \\
&= \max_{\|y\|_2 \leq Q} \mathcal{L}(\bar{x}^N, y) - \mathcal{L}(x^\star, \bar{y}^N) \\
&\leq \max_{\|y\|_2 \leq Q} \frac{1}{2\sum_{k=0}^N \rho_k} \left(\|(x^0, y^0) - (x^\star, y)\|_M^2 + \frac{\alpha L_f}{(1-\alpha\bar{\rho})\lambda_{\min}(M)} \|z^0 - z^\star\|_M^2 \right) \\
&= C_1 / (\sum_{k=1}^N \rho_k).
\end{aligned}$$

□

We need the following lemma to prove Theorem 2.

Lemma 3 (Davis (2015), Theorem 4.1)**.** *Suppose $\mathcal{T} : \mathbb{R}^n \to \mathbb{R}^n$ is an α-averaged operator with respect to $\|\cdot\|_M$, where $0 < \alpha < 1$ and $M \succ 0$. Let $z^\star \in \mathbf{Fix}\,\mathcal{T}$ and $z^0 \in \mathbb{R}^n$. For $\{\rho_k\} \subset (1, 1/\alpha)$, consider a sequence $\{z^k\}$ generated by the KM iteration:*

$$z^{k+1} = z^k + \rho_k(\mathcal{T}z^k - z^k).$$

If $\tau = \sup_{k \geq 0}(1 - \alpha\rho_k)\rho_k/\alpha > 0$, then we have

$$\|\mathcal{T}z^k - z^k\|_M^2 \leq \frac{\|z^0 - z^\star\|_M^2}{\tau(k+1)} \quad \textit{and} \quad \|\mathcal{T}z^k - z^k\|_M^2 = o\left(\frac{1}{k+1}\right). \tag{46}$$

Proof of Theorem 2. By condition (23), $\|z'\|_M^2 \geq \nu\|x'\|_2^2 + \epsilon\|y'\|_2^2$ for all $z' = (x', y')$. Then, in the same manner as the proof of Theorem 1, we put $z^- = z^k$, $z^+ = \tilde{z}^k$, $\rho = \rho_k$, $z_\rho = z^{k+1}$ in Lemma 1 and note that $1 < 1/\alpha < 2$ by the assumption $\nu > L_f/2$ to have

$$\begin{aligned}
2\rho_k(\mathcal{L}(\tilde{x}^k, y) - \mathcal{L}(x, \tilde{y}^k)) \leq \|z^k - z\|_M^2 - \|z^{k+1} - z\|_M^2 + \epsilon(1 - 2/\rho_k)\|y^k - y^{k+1}\|_2^2 \\
+ \left(\nu - \tfrac{2\nu - L_f}{\rho_k}\right)\|x^k - x^{k+1}\|_2^2.
\end{aligned} \tag{47}$$

The rest of the proof closely follows that of Davis (2015, Theorem 4.2). Note ν satisfies (15) and hence by Proposition 2, $R_{M^{-1}F}(I - M^{-1}G) : z^k \mapsto \tilde{z}^k$ is α-averaged with respect to $\|\cdot\|_M$. Let $z_\rho = (1-\rho)z^k + \rho\tilde{z}^k =: T_\rho z^k$ for any $\rho \in (0, 1/\alpha)$; for $\rho = \rho_k$, we have $z_\rho = z^{k+1}$. Then the map $T_\rho : z^k \mapsto z_\rho$ is $\alpha\rho$-averaged with respect to $\|\cdot\|_M$ and hence $\|z_\rho - z^\star\|_M \leq \|z^k - z^\star\|_M$. From Proposition 1 (1), we have $\|z_\rho - z^\star\|_M \leq \|z^0 - z^\star\|_M$, thus by the triangle inequality $\|z_\rho - z\|_M \leq \|z^0 - z^\star\|_M + \|z^\star - z\|_M$ for any $z \in \mathbb{R}^{p+l}$. Then we have

$$\begin{aligned}
(1/\rho)\langle z^k - z_\rho, z_\rho - z\rangle_M &= \langle \tilde{z}^k - z^k, z_\rho - z\rangle_M \\
&\leq \|\tilde{z}^k - z^k\|_M \|z_\rho - z\|_M \leq \frac{\|z^0 - z^\star\|_M}{\sqrt{\tau(k+1)}}(\|z^0 - z^\star\|_M + \|z^\star - z\|_M)
\end{aligned} \tag{48}$$

for all $\rho \in (0, 1/\alpha)$, where the last inequality is from Lemma 3.

Note that Lemma 1 (with the improvement (47) above) still holds if ρ_k is replaced by any $\rho \in (0, 1/\alpha)$ and z^{k+1} is replaced by z_ρ. Therefore we have

$$
\begin{aligned}
&\mathcal{L}(\tilde{x}^k, y) - \mathcal{L}(x, \tilde{y}^k) \\
&\quad \le \inf_{0<\rho<1/\alpha} \frac{1}{2\rho} \left(\|z^k - z\|_M^2 - \|z_\rho - z\|_M^2 - \epsilon(\tfrac{2}{\rho} - 1)\|y_\rho - y^k\|_2^2 + (\nu - \tfrac{2\nu - L_f}{\rho})\|x_\rho - x^k\|_2^2 \right) \\
&\quad = \inf_{0<\rho<1/\alpha} \frac{1}{2\rho} \left(2\langle z^k - z_\rho, z_\rho - z\rangle_M + \|z_\rho - z^k\|_M^2 - \epsilon(\tfrac{2}{\rho} - 1)\|y_\rho - y^k\|_2^2 + (\nu - \tfrac{2\nu - L_f}{\rho})\|x_\rho - x^k\|_2^2 \right) \\
&\quad \le \inf_{0<\rho<1/\alpha} \frac{1}{2\rho} \left(2\langle z^k - z_\rho, z_\rho - z\rangle_M + (\bar{\lambda} + \epsilon - \tfrac{2\epsilon}{\rho})\|y_\rho - y^k\|_2^2 + (\bar{\lambda} + \nu - \tfrac{2\nu - L_f}{\rho})\|x_\rho - x^k\|_2^2 \right) \\
&\quad = \inf_{0<\rho<1/\alpha} \frac{1}{\rho}\langle z^k - z_\rho, z_\rho - z\rangle_M + \frac{1}{2} \left((\bar{\lambda} + \epsilon - \tfrac{2\epsilon}{\rho})\|\tilde{y}^k - y^k\|_2^2 + (\bar{\lambda} + \nu - \tfrac{2\nu - L_f}{\rho})\|\tilde{x}^k - x^k\|_2^2 \right) \\
&\quad \le \frac{1}{\tilde{\rho}}\langle z^k - z_{\tilde{\rho}}, z_{\tilde{\rho}} - z\rangle_M
\end{aligned}
$$

by choosing a small $\tilde{\rho} \in (0, 1/\alpha)$ such that $\bar{\lambda} + \epsilon \le 2\epsilon/\tilde{\rho}$ and $\bar{\lambda} + \nu \le (2\nu - L_f)/\tilde{\rho}$, where $\bar{\lambda} = \lambda_{\max}(M)$. The first equality uses the cosine rule

$$2\langle a - b, c - b\rangle_M = -\|a - c\|_M^2 + \|a - b\|_M^2 + \|c - b\|_M^2$$

for any $a, b, c \in \mathbb{R}^{p+l}$. The desired result follows from (48).

The $o(1/\sqrt{k+1})$ rate is also from (48) and Lemma 3. □

We need the following lemma to prove Theorem 3.

Lemma 4. *Let $F : \mathbb{R}^{p+l} \to 2^{\mathbb{R}^{p+l}}$ is a maximally monotone operator, $G : \mathbb{R}^{p+l} \to \mathbb{R}^{p+l}$ is a single-valued operator and satisfies*

$$\langle Gz - Gz', z - z'\rangle \ge \|Gz - Gz'\|_{L^{-1}}^2, \quad \forall z, z' \in \mathbb{R}^{p+l},$$

where L is a symmetric positive definite matrix. If $M \succ (1/2)L$, then for a nondecreasing sequence $\{\beta_k\} \subset [0, \beta]$ such that there exists an $\eta > 0$ with

$$(1 - 3\beta_k - \eta)M \succeq \frac{(1 - \beta_k)^2}{2} L,$$

the sequence $\{z^k\}$ generated by (22) with $z^0 = z^{-1} \in \mathbb{R}^{p+l}$ satisfies the following.

1. *$\{z^k\}$ converges to a $z^\star \in \mathbf{zer}\,(F + G)$,*
2. *$\sum_{k=0}^{\infty} \|z^{k+1} - z^k\|_M^2 \le \frac{2(1-\beta_0)}{(1-\beta)\eta}\|z^0 - z^\star\|_M^2$, thus $\|z^{k+1} - z^k\|_M \to 0$,*

Proof of Theorem 3. From Lemma 1 with $z^- = \tilde{z}^k$, $z_\rho = z^+ = z^{k+1}$ and $\rho = 1$, we have

$$\begin{aligned}
2(\mathcal{L}(x^{k+1}, y) - \mathcal{L}(x, y^{k+1})) &\leq \|\tilde{z}^k - z\|_M^2 - \|z^{k+1} - z\|_M^2 - \|\tilde{z}^k - z^{k+1}\|_M^2 + L_f\|\tilde{x}^k - x^{k+1}\|_2^2 \\
&= 2\langle z^{k+1} - \tilde{z}^k, z - z^{k+1}\rangle_M + L_f\|\tilde{x}^k - x^{k+1}\|_2^2 \\
&= 2\langle z^{k+1} - z^k, z - z^{k+1}\rangle_M - 2\beta_k\langle z^k - z^{k-1}, z - z^{k+1}\rangle_M \\
&\quad + L_f\|x^{k+1} - x^k - \beta_k(x^k - x^{k-1})\|_2^2 \\
&\leq 2\langle z^{k+1} - z^k, z - z^{k+1}\rangle_M - 2\beta_k\langle z^k - z^{k-1}, z - z^k + z^k + z^{k+1}\rangle_M \\
&\quad + L_f\left((1+\beta_k)\|x^{k+1} - x^k\|_2^2 + \beta_k(1+\beta_k)\|x^k - x^{k-1}\|_2^2\right) \\
&\leq (\|z - z^k\|_M^2 - \|z - z^{k+1}\|_M^2 - \|z^{k+1} - z^k\|_M^2) \\
&\quad - \beta_k(\|z - z^{k-1}\|_M^2 - \|z - z^k\|_M^2 - \|z^k - z^{k-1}\|_M^2) \\
&\quad - 2\beta_k\langle z^k - z^{k-1}, z^k - z^{k+1}\rangle_M \\
&\quad + L_f\left((1+\beta_k)\|x^{k+1} - x^k\|_2^2 + \beta_k(1+\beta_k)\|x^k - x^{k-1}\|_2^2\right) \\
&\leq (\|z - z^k\|_M^2 - \|z - z^{k+1}\|_M^2) + \beta_k(\|z - z^k\|_M^2 - \|z - z^{k-1}\|_M^2) \\
&\quad + (\beta_k - 1)\|z^{k+1} - z^k\|_M^2 + 2\beta_k\|z^k - z^{k-1}\|_M^2 \\
&\quad + L_f\left((1+\beta_k)\|x^{k+1} - x^k\|_2^2 + \beta_k(1+\beta_k)\|x^k - x^{k-1}\|_2^2\right),
\end{aligned}$$

where the second and final inequalities use $2|\langle a, b\rangle_M| \leq \|a\|_M^2 + \|b\|_M^2$ (see Chambolle and Pock, 2016, Theorem 3). Because β_k is nondecreasing and $z^{-1} = z^0$, by summing the above inequality we have

$$\begin{aligned}
\sum_{k=0}^{N-1} 2(\mathcal{L}(x^{k+1}, y) - \mathcal{L}(x, y^{k+1})) &\leq (1 - \beta_0)\|z - z^0\|_M^2 + L_f(1+\bar{\beta})^2 \sum_{k=0}^{\infty}\|x^{k+1} - x^k\|_2^2 \\
&\quad + \beta_{N-1}\|z - z^{N-1}\|_M^2 - \|z - z^N\|_M^2 + (\beta_{N-1} - 1)\|z^N - z^{N-1}\|_M^2.
\end{aligned}$$

The second line is nonpositive because

$$\begin{aligned}
&\beta_{N-1}\|z - z^{N-1}\|_M^2 - \|z - z^N\|_M^2 + (\beta_{N-1} - 1)\|z^N - z^{N-1}\|_M^2 \\
&\quad = \beta_{N-1}\|z - z^N + z^N - z^{N-1}\|_M^2 - \|z - z^N\|_M^2 + (\beta_{N-1} - 1)\|z^N - z^{N-1}\|_M^2 \\
&\quad \leq 2\beta_{N-1}(\|z - z^N\|_M^2 + \|z^N - z^{N-1}\|_M^2) - \|z - z^N\|_M^2 + (\beta_{N-1} - 1)\|z^N - z^{N-1}\|_M^2 \\
&\quad = (2\beta_{N-1} - 1)\|z - z^N\|_M^2 + (3\beta_{N-1} - 1)\|z^N - z^{N-1}\|_2^2
\end{aligned}$$

and $\beta_k \leq \bar{\beta} \leq 1/3$ for all k. Note the conditions for Lemma 4 are met and thus by Result 2 of Proposition 1 we have

$$\sum_{k=0}^{\infty}\|x^{k+1} - x^k\|_2^2 \leq \frac{1}{\lambda_{\min}(M)}\sum_{k=0}^{\infty}\|z^{k+1} - z^k\|_M^2 \leq \frac{2(1-\beta_0)}{(1-\bar{\beta})\eta\lambda_{\min}(M)}\sum_{k=0}^{\infty}\|z^0 - z^\star\|_M^2.$$

Finally, the desired result follows from the convexity-concavity of $\mathcal{L}(x, y)$:

$$\mathcal{L}(\bar{x}^N, y) - \mathcal{L}(x, \bar{y}^N) \leq \frac{1}{N}\sum_{k=0}^{N-1}(\mathcal{L}(x^{k+1}, y) - \mathcal{L}(x, y^{k+1})).$$

□